



## INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<b>(51) International Patent Classification <sup>6</sup> :</b> <b>G06F 17/00</b>	<b>A2</b>	<b>(11) International Publication Number:</b> <b>WO 99/28832</b> <b>(43) International Publication Date:</b> 10 June 1999 (10.06.99)
<b>(21) International Application Number:</b> PCT/US98/25535 <b>(22) International Filing Date:</b> 3 December 1998 (03.12.98)  <b>(30) Priority Data:</b> 09/985,114 4 December 1997 (04.12.97) US  <b>(71) Applicant:</b> MICROSOFT CORPORATION [US/US]; One Microsoft Way, Redmond, WA 98052 (US).  <b>(72) Inventors:</b> THIESSON, Bo; 1502 3rd Place, Kirkland, WA 98033 (US). MEEK, Christopher, A.; 12935 N.E. 71st Street, Kirkland, WA 98033 (US). CHICKERING, David, Maxwell; 7573 Old Redmond Road #12, Redmond, WA 98052 (US). HECKERMAN, David, Earl; 648 West LK Sammamish Lane N.E., Bellevue, WA 98008 (US).  <b>(74) Agent:</b> MICHAELSON, Peter, L.; Michaelson & Wallace, Parkway 109 Office Center, 328 Newman Springs Road, P.O. Box 8489, Red Bank, NJ 07701 (US).		<b>(81) Designated States:</b> CA, CN, JP, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).  <b>Published</b> <i>Without international search report and to be republished upon receipt of that report.</i>
<b>(54) Title:</b> MIXTURES OF BAYESIAN NETWORKS  <b>(57) Abstract</b>  One aspect of the invention is the construction of mixtures of Bayesian networks. Another aspect of the invention is the use of such mixtures of Bayesian networks to perform inferencing. A mixture of Bayesian networks (MBN) consists of plural hypothesis-specific Bayesian networks (HSBNs) having possibly hidden and observed variables. A common external hidden variable is associated with the MBN, but is not included in any of the HSBNS. The number of HSBNS in the MBN corresponds to the number of states of the common external hidden variable, and each HSBNS is based upon the hypothesis that the common external hidden variable is in a corresponding one of those states. In one mode of the invention, the MBN having the highest MBN score is selected for use in performing inferencing. In another mode of the invention, some or all of the MBNs are retained as a collection of MBNs which perform inferencing in parallel, their outputs being weighted in accordance with the corresponding MBN scores and the MBN collection output being the weighted sum of all the MBN outputs. In one application of the invention, collaborative filtering may be performed by defining the observed variables to be choices made among a sample of users and the hidden variables to be the preferences of those users.		

**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

-1-

**MIXTURES OF BAYESIAN NETWORKS**Field of the Invention

The present invention relates generally to data processing systems and, more particularly, to the generation of Bayesian networks.

Background of the Invention

The advent of artificial intelligence within computer science has brought an abundance of decision-support systems. Decision-support systems are computer systems in which decisions, typically rendered by humans, are recommended and sometimes made. In creating decision-support systems, computer scientists seek to provide decisions with the greatest possible accuracy. Thus, computer scientists strive to create decision-support systems that are equivalent to or more accurate than a human expert. Applications of decision-support systems include medical diagnosis, troubleshooting computer networks, or other systems wherein a decision is based upon identifiable criteria.

One of the most promising new areas for research in decision-support systems is Bayesian networks. A Bayesian network is a representation of the probabilistic relationships among distinctions about the world. Each distinction, sometimes called a

-2-

variable, can take on one of a mutually exclusive and exhaustive set of possible states. A Bayesian network is expressed as an acyclic-directed graph where the variables correspond to nodes and the relationships between the nodes correspond to arcs. Figure 1 depicts an exemplary Bayesian network 101. In Figure 1 there are three variables,  $X_1$ ,  $X_2$ , and  $X_3$ , which are represented by nodes 102, 106 and 110, respectively. This Bayesian network contains two arcs 104 and 108. Associated with each variable in a Bayesian network is a set of probability distributions. Using conditional probability notation, the set of probability distributions for a variable can be denoted by  $p(x_i | \Pi_i, \xi)$ , where "p" refers to the probability distribution, where " $\Pi_i$ " denotes the parents of variable  $X_i$  and where " $\xi$ " denotes the knowledge of the expert. The Greek letter " $\xi$ " indicates that the Bayesian network reflects the knowledge of an expert in a given field. Thus, this expression reads as follows: the probability distribution for variable  $X_i$  given the parents of  $X_i$  and the knowledge of the expert. For example,  $X_1$  is the parent of  $X_2$ . The probability distributions specify the strength of the relationships between variables. For instance, if  $X_1$  has two states (true and false), then associated with  $X_1$  is a single probability distribution  $p(x_1 | \xi)$  and associated with  $X_2$  are two probability distributions  $p(x_2 | x_1 = t, \xi)$  and

-3-

$p(x_2|x_1 = f, \xi)$ . In the remainder of this specification,  $\xi$  is not specifically mentioned.

The arcs in a Bayesian network convey  
5 dependence between nodes. When there is an arc between  
two nodes, the probability distribution of the first  
node depends upon the value of the second node when the  
direction of the arc points from the second node to the  
first node. For example, node 106 depends upon  
10 node 102. Therefore, nodes 102 and 106 are said to be  
conditionally dependent. Missing arcs in a Bayesian  
network convey conditional independencies. For  
example, node 102 and node 110 are conditionally  
independent given node 106. However, two variables  
15 indirectly connected through intermediate variables are  
conditionally dependent given lack of knowledge of the  
values ("states") of the intermediate variables.  
Therefore, if the value for node 106 is known, node 102  
and node 110 are conditionally dependent.

20

In other words, sets of variables X and Y are  
said to be conditionally independent, given a set of  
variables Z, if the probability distribution for X  
given Z does not depend on Y. If Z is empty, however,  
25 X and Y are said to be "independent" as opposed to  
conditionally independent. If X and Y are not  
conditionally independent, given Z, then X and Y are  
said to be conditionally dependent given Z.

-4-

The variables used for each node may be of different types. Specifically, variables may be of two types: discrete or continuous. A discrete variable is a variable that has a finite or countable number of states, whereas a continuous variable is a variable that has an uncountably infinite number of states. All discrete variables considered in this specification have a finite number of states. An example of a discrete variable is a Boolean variable. Such a variable can assume only one of two states: "true" or "false." An example of a continuous variable is a variable that may assume any real value between -1 and 1. Discrete variables have an associated probability distribution. Continuous variables, however, have an associated probability density function ("density"). Where an event is a set of possible outcomes, the density  $p(x)$  for a variable "x" and events "a" and "b" is defined as:

$$p(x) = \lim_{a \rightarrow b} \left[ \frac{p(a \leq x \leq b)}{|a - b|} \right]$$

where  $p(a \leq x \leq b)$  is the probability that x lies between a and b. Conventional systems for generating Bayesian networks cannot use continuous variables in their nodes.

Figure 2 depicts an example Bayesian network for troubleshooting automobile problems. The Bayesian network of Figure 2 contains many variables 202, 204,

-5-

206, 208, 210, 212, 214, 216, 218, 220, 222, 224, 226, 228, 230, 232, and 234, relating to whether an automobile will work properly, and arcs 236, 238, 240, 242, 244, 246, 248, 250, 252, 254, 256, 258, 260, 262, 264, 268. A few examples of the relationships between the variables follow. For the radio 214 to work properly, there must be battery power 212 (arc 246). Battery power 212, in turn, depends upon the battery working properly 208 and a charge 210 (arcs 242 and 244). The battery working properly 208 depends upon the battery age 202 (arc 236). The charge 210 of the battery depends upon the alternator 204 working properly (arc 238) and the fan belt 206 being intact (arc 240). The battery age variable 202, whose values lie from zero to infinity, is an example of a continuous variable that can contain an infinite number of values. However, the battery variable 208 reflecting the correct operations of the battery is a discrete variable being either true or false.

20

The automobile troubleshooting Bayesian network also provides a number of examples of conditional independence and conditional dependence. The nodes operation of the lights 216 and battery power 212 are dependent, and the nodes operation of the lights 216 and operation of the radio 214 are conditionally independent given battery power 212. However, the operation of the radio 214 and the operation of the lights 216 are conditionally dependent. The concept of conditional dependence and

30

-6-

conditional independence can be expressed using conditional probability notation. For example, the operation of the lights 216 is conditionally dependent on battery power 212 and conditionally independent of the radio 214 given the battery power 212. Therefore, the probability of the lights working properly 216 given both the battery power 212 and the radio 214 is equivalent to the probability of the lights working properly given the battery power alone,

$P(\text{Lights}|\text{Battery Power}, \text{Radio}) = P(\text{Lights}|\text{Battery Power})$ . An example of a conditional dependence relationship is the probability of the lights working properly 216 given the battery power 212 which is not equivalent to the probability of the lights working properly given no information. That is,

$p(\text{Lights}|\text{Battery Power}) \neq p(\text{Lights})$ .

There are two conventional approaches for constructing Bayesian networks. Using the first approach ("the knowledge-based approach"), a person known as a knowledge engineer interviews an expert in a given field to obtain the knowledge of the expert about the field of expertise of the expert. The knowledge engineer and expert first determine the distinctions of the world that are important for decision making in the field of the expert. These distinctions correspond to the variables of the domain of the Bayesian network. The "domain" of a Bayesian network is the set of all variables in the Bayesian network. The knowledge engineer and the expert next determine the dependencies



-7-

among the variables (the arcs) and the probability distributions that quantify the strengths of the dependencies.

5                   In the second approach ("called the data-based approach"), the knowledge engineer and the expert first determine the variables of the domain. Next, data is accumulated for those variables, and an algorithm is applied that creates a Bayesian network  
10                   from this data. The accumulated data comes from real world instances of the domain. That is, real world instances of decision making in a given field. Conventionally, this second approach exists for domains containing only discrete variables.

15                   After the Bayesian network has been created, the Bayesian network becomes the engine for a decision-support system. The Bayesian network is converted into a computer-readable form, such as a file  
20                   and input into a computer system. Then, the computer system uses the Bayesian network to determine the probabilities of variable states given observations, determine the benefits of performing tests, and ultimately recommend or render a decision. Consider an  
25                   example where a decision-support system uses the Bayesian network of Figure 2 to troubleshoot automobile problems. If the engine for an automobile did not start, the decision-based system could request an observation of whether there was gas 224, whether the  
30                   fuel pump 226 was in working order by possibly

-8-

performing a test, whether the fuel line 228 was obstructed, whether the distributor 230 was working, and whether the spark plugs 232 were working. While the observations and tests are being performed, the Bayesian network assists in determining which variable should be observed next.

U.S. application Serial No. \_\_\_\_\_  
filed \_\_\_\_\_ entitled "Generating Improved Belief  
10 Networks" describes an improved system and method for generating Bayesian networks (also known as "belief networks") that utilize both expert data received from an expert ("expert knowledge") and data received from real world instances of decisions made ("empirical  
15 data"). By utilizing both expert knowledge and empirical data, the network generator provides an improved Bayesian network that is more accurate than conventional Bayesian networks. In addition, the exemplary embodiment facilitates the use of continuous  
20 variables in Bayesian networks and handles missing data in the empirical data that is used to construct Bayesian networks.

Expert knowledge consists of two components:  
25 an equivalent sample size or sizes ("sample size"), and the prior probabilities of all possible Bayesian-network structures ("priors on structures"). The effective sample size is the effective number of times that the expert has rendered a specific decision.  
30 For example, a doctor with 20 years of experience

-9-

diagnosing a specific illness may have an effective sample size in the hundreds. The priors on structures refers to the confidence of the expert that there is a relationship between variables (e.g., the expert is 70 percent sure that two variables are related). The priors on structures can be decomposed for each variable-parent pair known as the "prior probability" of the variable-parent pair.

Empirical data is typically stored in a database. An example of acquiring empirical data can be given relative to the Bayesian network of Figure 2. If, at a service station, a log is maintained for all automobiles brought in for repair, the log constitutes empirical data. The log entry for each automobile may contain a list of the observed state of some or all of the variables in the Bayesian network. Each log entry constitutes a case. When one or more variables are unobserved in a case, the case containing the unobserved variable is said to have "missing data." Thus, missing data refers to when there are cases in the empirical data database that contain no observed value for one or more of the variables in the domain. An assignment of one state to each variable in a set of variables is called an "instance" of that set of variables. Thus, a "case" is an instance of the domain. The "database" is the collection of all cases.

An example of a case can more clearly be described relative to the Bayesian network of Figure 2.

-10-

A case may consist of the battery age 202 being 2.132 years old, the battery working properly 208 being true, the alternator working properly 204 being true, the fan belt being intact 206 being true, the charge 210 being sufficient, the battery power 212 being sufficient, the starter working properly 220 being true, the engine turning over 218 being true, the amount of gas 224 being equal to 5.3 gallons, the fuel pump working properly 226 being true, the fuel line working properly 228 being true, the distributor working properly 230 being false, the spark plugs working properly 232 being true and the engine starting 234 being false. In addition, the variables for the gas gauge 222, the radio working properly 214 and the lights working properly 216 may be unobserved. Thus, the above-described case contains missing data.

#### Background Relative to Decision Graphs:

Although Bayesian networks are quite useful in decision-support systems, Bayesian networks require a significant amount of storage. For example, in the Bayesian network 300 of Figure 3A, the value of nodes X and Y causally influences the value of node Z. In this example, nodes X, Y, and Z have binary values of either 0 or 1. As such, node Z maintains a set of four probabilities, one probability for each combination of the values of X and Y, and stores these probabilities into a table 320 as shown in Figure 3B. When performing probabilistic inference, it is the

-11-

probabilities in table 320 that are accessed. As can be seen from table 320, only the probabilities for Z equaling 0 are stored; the probabilities for Z equaling 1 need not be stored as they are easily derived by subtracting the probability of when Z equals 0 from 1. As the number of parents of a node increases, the table in the node that stores the probabilities becomes multiplicatively large and requires a significant amount of storage. For example, a node having binary values with 10 parents that also have binary values requires a table consisting of 1,024 entries. And, if either the node or one of its parents has more values than a binary variable, the number of probabilities in the table increases multiplicatively.

To improve the storage of probabilities in a Bayesian network node, some conventional systems use a tree data structure. A tree data structure is an acyclic, undirected graph where each vertex is connected to each other vertex via a single path. The graph is acyclic in that there is no path that both emanates from a vertex and returns to the same vertex, where each edge in the path is traversed only once. Figure 3C depicts an example tree data structure 330 that stores into its leaf vertices 336-342 the probabilities shown in table 320 of Figure 3B. Assuming that a decision-support system performs probabilistic inference with X's value being 0 and Y's value being 1, the following steps occur to access the appropriate probability in the tree data structure 330:

-12-

First, the root vertex 332, vertex X, is accessed, and its value determines the edge or branch to be traversed. In this example, X's value is 0, so edge 344 is traversed to vertex 334, which is vertex Y. Second, after reaching vertex Y, the value for this vertex determines which edge is traversed to the next vertex. In this example, the value for vertex Y is 1, so edge 346 is traversed to vertex 338, which is a leaf vertex. Finally, after reaching the leaf vertex 338, which stores the probability for Z equaling 0 when  $X = 0$  and  $Y = 1$ , the appropriate probability can be accessed.

As compared to a table, a tree is a more efficient way of storing probabilities in a node of a Bayesian network, because it requires less space. However, tree data structures are inflexible in the sense that they can not adequately represent relationships between probabilities. For example, because of the acyclic nature of tree data structures, a tree cannot be used to indicate some types of equality relationships where multiple combinations of the values of the parent vertices have the same probability (i.e., refer to the same leaf vertex). This inflexibility requires that multiple vertices must sometimes store the same probabilities, which is wasteful. It is thus desirable to improve Bayesian networks with tree distributions.

-13-

Background Relative to Collaborative Filtering:

Collaborative filtering systems have been developed that predict the preferences of a user. The term "collaborative filtering" refers to predicting the preferences of a user based on known attributes of the user, as well as known attributes of other users. For example, a preference of a user may be whether they would like to watch the television show "I Love Lucy" and the attributes of the user may include their age, gender, and income. In addition, the attributes may contain one or more of the user's known preferences, such as their dislike of another television show. A user's preference can be predicted based on the similarity of that user's attributes to other users. For example, if all users over the age of 50 with a known preference happen to like "I Love Lucy" and if that user is also over 50, then that user may be predicted to also like "I Love Lucy" with a high degree of confidence. One conventional collaborative filtering system has been developed that receives a database as input. The database contains attribute-value pairs for a number of users. An attribute is a variable or distinction, such as a user's age, gender or income, for predicting user preferences. A value is an instance of the variable. For example, the attribute age may have a value of 23. Each preference contains a numeric value indicating whether the user likes or dislikes the preference (e.g., 0 for dislike and 1 for like). The data in the

-14-

database is obtained by collecting attributes of the users and their preferences.

It should be noted that conventional collaborative filtering systems can typically only utilize numerical attributes. As such, the values for non-numerical attributes, such as gender, are transposed into a numerical value, which sometimes reduces the accuracy of the system. For example, when a variable has three non-numerical states, such as vanilla, chocolate and strawberry, transposing these states into a numerical value will unintentionally indicate dissimilarity between the states. That is, if vanilla were assigned a value of 1, chocolate 2 and strawberry 3, the difference between each value indicates to the system how similar each state is to each other. Therefore, the system may make predictions based on chocolate being more similar to both vanilla and strawberry than vanilla is similar to strawberry. Such predictions may be based on a misinterpretation of the data and lead to a reduction in the accuracy of the system.

In performing collaborative filtering, the conventional system first computes the correlation of attributes between a given user "v" and each other user "u" (except v) in the database. The computation of the "correlation" is a well-known computation in the field of statistics. After computing the correlation, the conventional system computes, for example, the



-15-

preference of a user "v" for a title of a television show "t" as follows:

$$\text{pref}(t, v) = \langle \text{pref}(t) \rangle + \frac{\sum_u (\text{pref}(t, u) - \langle \text{pref}(t) \rangle) \text{corr}(u, v)}{\sum_u \text{corr}(u, v)}$$

5

where "pref(t, v)" is the preference of user "v" for title "t," where "< pref(t) >" is the average preference of title "t" by all users, where "pref(t, u)" is the preference of user "u" for title "t," where "corr(u, v)" is the correlation of users "u" and "v," and the sums run over the users "u" that have expressed a preference for title "t." One drawback to this conventional system is that the entire database must be examined when predicting preferences, which requires a significant amount of processing time.

15

One way to improve upon this conventional system is to utilize a clustering algorithm. Using this approach, a collaborative filtering system uses any of a number of well-known clustering algorithms to divide the database into a number of clusters. For example, the algorithms described in KoJain, Algorithms for Clustering Data (1988) can be used. Each cluster contains the data of users whose preferences tend to be similar. As such, when predicting the preferences of one user in a cluster, only the preferences of the other users in the cluster need to be examined and not the preferences of all other users in the database. A

20

25

-16-

collaborative filtering system that utilizes a clustering algorithm receives as input a database, as described above, a guess of the number of clusters and a distance metric. The guess of the number of clusters is provided by an administrator of the collaborative filtering system based on their own knowledge of how many clusters the database can probably be divided into. The distance metric is a metric provided by the administrator for each user in the database that estimates how similar one user is to each other in the database based on user's preferences and attributes. The distance metric is a range between 0 and 1 with 0 indicating that two users are least similar and 1 indicating that two users are most similar. This similarity is expressed as a numerical value. Each user will have a distance metric for every other user. Thus, the distance metrics are conveniently represented by an N-by-N matrix, where "N" is the number of users. After receiving the number of clusters and the distance metric, the clustering algorithm identifies the clusters.

The clustering algorithm outputs a list of the users in the database and a cluster number assigned to each user. To determine the preferences of a user, the other users within that user's cluster are examined. For example, if the system is attempting to determine whether a user would like the television show "I Love Lucy," the other users within that cluster are examined. If there are six other users within the

-17-

cluster and five out of the six like "I Love Lucy," then it is likely that so will the user.

Although utilizing a clustering algorithm may be an improvement over the previously-described conventional system, it has limitations. One such limitation is that the exact number of clusters is determined manually, which renders the algorithm prone to human error. Another limitation is that all attributes are numerical and as such, the values of non-numerical attributes must be transposed into numerical values. Based upon the above-described limitations of conventional collaborative filtering systems, it is desirable to improve collaborative filtering systems.

#### Summary of the Invention

One aspect of the invention is the construction of mixtures of Bayesian networks. Another aspect of the invention is the use of such mixtures of Bayesian networks to perform inferencing. A mixture of Bayesian networks (MBN) consists of plural hypothesis-specific Bayesian networks (HSBNs) having possibly hidden and observed variables. A common external hidden variable is associated with the MBN, but is not included in any of the HSBNS. The number of HSBNS in the MBN corresponds to the number of states of the common external hidden variable, and each HSBN models the world under the hypothesis that the common

-18-

external hidden variable is in a corresponding one of those states.

5       The MBN structure is initialized as a  
collection of identical HSBNs whose discrete hidden  
variables are connected to all observed variables and  
whose continuous hidden variables are connected only to  
each of the continuous observed variables, the  
directionality being from hidden variable to observed  
10       variable.

      In constructing the MBN, the parameters of the  
current HSBNs are improved using an  
expectation-maximization process applied for training  
15       data. The expectation-maximization process is iterated  
to improve the network performance in predicting the  
training data, until some criteria has been met. Early  
in the process, this criteria may be a fix number of  
iterations which may itself be a function of the number  
20       of times the overall learning process has iterated.  
Later in the process, this criteria may be convergence  
of the parameters to a near optimum network performance  
level.

25       Then, expected complete-model sufficient  
statistics are generated from the training data. The  
expected complete-model sufficient statistics are  
generated as follows: first, a vector is formed for  
each observed case in the training data. Each entry in  
30       the vector corresponds to a configuration of the

-19-

discrete variables. Each entry is itself a vector with subentries. The subentries for a given case are (1) the probability that, given the data of the particular case, the discrete variables are in the configuration  
5 corresponding to the entry's position within the vector, and (2) information defining the state of the continuous variables in that case multiplied by the probability in (1). These probabilities are computed by conventional techniques using the MBN in its current form. In this  
10 computation, conditional probabilities derived from the individual HSBNs are weighted and then summed together. The individual weights correspond to the current probabilities of the common external hidden variable being in a corresponding one of its states. These  
15 weights are computed from the MBN in its current form using conventional techniques. Once such vectors are formed for all the cases represented by the training data, the expected complete-model sufficient statistics are then generated by summing the vectors together,  
20 i.e., summing the vectors over all cases.

After computation of the expected  
complete-model sufficient statistics for the MBN, the structures of the HSBNs are searched for changes which  
25 improve the HSBN's score or performance in predicting the training data given the current parameters. The MBN score preferably is determined by the HSBN scores, the score for the common hidden external variable, and a correction factor. If the structure of any HSBN changes  
30 as a result of this fast search, the prior steps

-20-

beginning with the expectation-maximization process are repeated. The foregoing is iteratively repeated until the network structure stabilizes. At this point the current forms of the HSBNs are saved as the MBN. An MBN  
5 is thus generated for each possible combination of number of states of the hidden discrete variables including the common external hidden variable, so that a number of MBNs is produced in accordance with the number of combinations of numbers of states of the hidden  
10 discrete variables.

In one mode of the invention, the MBN having the highest MBN score is selected for use in performing inferencing. In another mode of the invention, some or  
15 all of the MBNs are retained as a collection of MBNs which perform inferencing in parallel, their outputs being weighted in accordance with the corresponding MBN scores and the MBN collection output being the weighted sum of all the MBN outputs.

20 Collaborative filtering may be performed by defining the observed variables to be the preferences of those users. The common hidden discrete variable then may be an unknown class variable, which is never  
25 discovered in the network generation process nor during the use of the MBN to perform inferencing.

-21-

Brief Description of the Drawings

Figure 1 depicts an example of a conventional Bayesian network.

5

Figure 2 depicts an example conventional Bayesian network for troubleshooting automobile problems.

10

Figure 3A depicts a conventional Bayesian network.

Figure 3B depicts a table containing the probabilities for one of the nodes of the conventional Bayesian network of Figure 3A.

15

Figure 3C depicts a tree data structure containing the probabilities for one of the nodes of the Bayesian network of Figure 3A.

20

Figure 4 depicts a computer system suitable for practicing an exemplary embodiment of the present invention.

25

Figure 5 depicts a functional overview of the Bayesian network generator of the exemplary embodiment.

Figure 6 depicts the Bayesian network generator of an exemplary embodiment in a computer system suitable for practicing the present invention.

30

-22-

Figure 7 depicts an exemplary Bayesian network consisting of a mixture of Bayesian networks in accordance with the invention.

5                   Figure 8 depicts one exemplary hypothesis-specific Bayesian network in the mixture of Bayesian networks of Figure 7.

10                   Figure 9 depicts another exemplary hypothesis-specific Bayesian network in the mixture of Bayesian networks of Figure 7.

15                   Figure 10 depicts an initial Bayesian network.

                  Figure 11 depicts a mixture of hypothesis specific networks corresponding to the network of Figure 10.

20                   Figure 12 illustrates a method of generating mixtures of Bayesian networks in accordance with a first exemplary embodiment of the invention.

25                   Figure 13 illustrates a method of generating mixtures of Bayesian networks in accordance with a second exemplary embodiment of the invention.

30                   Figure 14 illustrates a method of generating mixtures of Bayesian networks in accordance with a third exemplary embodiment of the invention.



-23-

Figure 15 illustrates a method of generating mixtures of Bayesian networks in accordance with a fourth exemplary embodiment of the invention.

5                   Figure 16 illustrates an inferencing apparatus including a mixture of Bayesian networks in accordance with one aspect of the invention.

10                   Figure 17 illustrates an inferencing apparatus including a collection of mixtures of Bayesian networks in accordance with another aspect of the invention.

15                   Figure 18 depicts a more detailed diagram of the Bayesian network generator of Figure 6.

20                   Figure 19 depicts a high-level flow chart of the steps performed by the scoring mechanism of Figure 18.

                  Figure 20 depicts a flow chart of the steps performed by the calculate discrete score process of Figure 19.

25                   Figures 21 and 21B depict a flow chart of the steps performed by the calculate continuous score process of Figure 19.

-24-

Figure 22 depicts a flow chart of the steps performed by the calculate mixed score process of Figure 19.

5            Figures 23A and 23B depict a flow chart of the steps performed by the network adjuster of Figure 18.

10           Figure 24 depicts a decision graph data structure as used by the Bayesian network of an exemplary embodiment of the present invention.

15           Figure 25A depicts a Bayesian network of an exemplary embodiment of the present invention.

            Figure 25B depicts a decision graph suitable for use in one of the nodes of the Bayesian network of Figure 25A.

20           Figure 25C depicts a Bayesian network of an alternative embodiment of the present invention, which contains cycles.

25           Figure 26A depicts a flowchart of the steps performed by one implementation of the Bayesian network generator depicted in Figure 6.

30           Figure 26B depicts a flowchart of the steps performed by the Bayesian network generator when generating candidate decision graphs.

-25-

Figure 27A depicts an exemplary decision graph.

5 Figure 27B depicts the exemplary decision graph of Figure 28A after a complete split has been performed on one of the leaf nodes.

10 Figure 27C depicts the exemplary decision graph of Figure 28A after a binary split has been performed on one of the leaf nodes.

15 Figure 27D depicts the exemplary decision graph of Figure 28A after a merge has been performed on two of the leaf nodes of the decision graph of Figure 28A.

20 Figure 28 depicts a flowchart of the steps performed by the web analyzer of an exemplary embodiment of the present invention.

Figure 29 depicts a hypothesis-specific Bayesian network in an example relative to collaborative filtering.

25 Detailed Description of the Invention

Exemplary Operating Environment:

30 Figure 4 and the following discussion are intended to provide a brief, general description of a

-26-

suitable computing environment in which the invention may be implemented. Although not required, the invention will be described in the general context of computer-executable instructions, such as program modules, being executed by a personal computer.

Generally, program modules include processes, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Moreover, those skilled in the art will appreciate that the invention may be practiced with other computer system configurations, including hand-held devices, multiprocessor systems, microprocessor-based or programmable consumer electronics, network PCs, minicomputers, mainframe computers, and the like. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located both local and remote memory storage devices.

With reference to Figure 4, an exemplary system for implementing the invention includes a general purpose computing device in the form of a conventional personal computer 420, including a processing unit 421, a system memory 422, and a system bus 423 that couples various system components including the system memory to the processing unit 421. The system bus 423 may be any of several types of bus

-27-

structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. The system memory includes read only memory (ROM) 424 and random access memory (RAM) 425. A basic input/output system 426 (BIOS), containing the basic process that helps to transfer information between elements within the personal computer 420, such as during start-up, is stored in ROM 424. The personal computer 420 further includes a hard disk drive 427 for reading from and writing to a hard disk, not shown, a magnetic disk drive 428 for reading from or writing to a removable magnetic disk 429, and an optical disk drive 430 for reading from or writing to a removable optical disk 431 such as a CD ROM or other optical media. The hard disk drive 427, magnetic disk drive 428, and optical disk drive 430 are connected to the system bus 423 by a hard disk drive interface 432, a magnetic disk drive interface 433, and an optical drive interface 434, respectively. The drives and their associated computer-readable media provide nonvolatile storage of computer readable instructions, data structures, program modules and other data for the personal computer 420. Although the exemplary environment described herein employs a hard disk, a removable magnetic disk 429 and a removable optical disk 431, it should be appreciated by those skilled in the art that other types of computer readable media which can store data that is accessible by a computer, such as magnetic cassettes, flash memory cards, digital video disks,

-28-

Bernoulli cartridges, random access memories (RAMs), read only memories (ROM), and the like, may also be used in the exemplary operating environment.

5                   A number of program modules may be stored on the hard disk, magnetic disk 429, optical disk 431, ROM 424 or RAM 425, including an operating system 435, one or more application programs 436, other program modules 437, and program data 438. A user may enter  
10                   commands and information into the personal computer 420 through input devices such as a keyboard 440 and pointing device 442. Other input devices (not shown) may include a microphone, joystick, game pad, satellite dish, scanner, or the like. These and other input  
15                   devices are often connected to the processing unit 421 through a serial port interface 446 that is coupled to the system bus, but may be connected by other interfaces, such as a parallel port, game port or a universal serial bus (USB). A monitor 447 or other  
20                   type of display device is also connected to the system bus 423 via an interface, such as a video adapter 448. In addition to the monitor, personal computers typically include other peripheral output devices (not shown), such as speakers and printers.

25                   The personal computer 420 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 449. The remote computer 449 may be another  
30                   personal computer, a server, a router, a network PC, a

-29-

peer device or other common network node, and typically includes many or all of the elements described above relative to the personal computer 420, although only a memory storage device 450 has been illustrated in Figure 4. The logical connections depicted in Figure 4 include a local area network (LAN) 451 and a wide area network (WAN) 452. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and Internet.

10

When used in a LAN networking environment, the personal computer 420 is connected to the local network 451 through a network interface or adapter 453. When used in a WAN networking environment, the personal computer 420 typically includes a modem 454 or other means for establishing communications over the wide area network 452, such as the Internet. The modem 454, which may be internal or external, is connected to the system bus 423 via the serial port interface 446. In a networked environment, program modules depicted relative to the personal computer 420, or portions thereof, may be stored in the remote memory storage device. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

20

25

-30-

Introduction to Mixtures of Bayesian Networks:

Figure 5 depicts a functional overview of the MBN generator of an exemplary embodiment. In order to use the MBN generator of the exemplary embodiment, a knowledge engineer first obtains expert knowledge from an expert in a given field (step 402). This expert knowledge includes one or more sample sizes and structure priors which includes the expert's prior probability that  $C$  has  $|C|$  states,  $p(|C|)$  and the expert's prior probability for each HSBN structure given  $c$  and  $p(B_c^e | |C|)$ . The knowledge engineer then obtains empirical data from real world invocations of decision making in the given field (step 404). After obtaining the expert knowledge and the empirical data, the knowledge engineer invokes the network generator of the exemplary embodiment to create an improved MBN that can then be used as the basis for a decision-support system (step 406). Although step 402 has been described as occurring before step 404, one skilled in the art will appreciate that step 404 may occur before step 402.

Figure 6 depicts the MBN generator of an exemplary embodiment in a computer system of the type depicted in Figure 4 suitable for practicing the exemplary embodiment of the present invention. The MBN generator of the exemplary embodiment 502 resides within a memory 304 and receives empirical data 504 and expert knowledge 506 as input. The expert



-31-

knowledge 506 typically comprises a sample size, and the priors on structures. Both the empirical data 504 and the expert knowledge 506 reside in a permanent storage device 306. The empirical data 504 is typically comprised of cases stored in a database ("the empirical data database"). In response to receiving both the empirical data 504 and the expert knowledge 506, the MBN generator 502 of the exemplary embodiment generates an MBN 508. The memory 304 and permanent storage 306 are connected to a central processing unit 302, a display 308 which may be a video display, and an input device 310.

Two types of problems that are addressed by the present invention are prediction tasks and clustering tasks.

A database of observed cases over a set of variables is given. The prediction problem is to learn the statistical relationships among those variables for prediction. The clustering problem is to group the rows of the database into groups so that groups of similar users can be discovered and properties of the groups can be presented. The invention provides a flexible and rich class of models (for both of these problems) and provide algorithms to learn which model from this class of models best fits the data. The class of models employed by the invention is called a mixture of Bayesian networks (MBN). The processes for learning MBNs include several advantageous features including:

-32-

- (a) interleaving parameter and structural search,
- (b) expected complete model sufficient statistics, and
- (c) an outer loop for determining the number of states of the discrete hidden variables.

5

The present invention is embodied in a mixture of Bayesian networks, which corresponds to a graphical model as shown in Figure 7. C, the class variable, is a discrete variable, that is not observed, 10 O is a set of observed variables and H is a set of unobserved (hidden) variables. As one example, C can have two possible values. In this case, the conditional distribution of sets of variables O and H given C = 0 might be represented by the Bayesian 15 network in Figure 8 and the conditional distribution of sets of variables O and H given C = 1 in Figure 9. Both sets of variables O and H may contain a combination of discrete and continuous variables. The only restriction is that no continuous variable can 20 point at a discrete variable in any of the Bayesian networks. Given a database of observations for the variables in O, the goal is to select the number of values for the class C (i.e., |C|), the parameters  $\theta_c$  (that describe the percentage of the database 25 attributed to the c'th Bayesian network), and the |C| Bayesian network structures and their parameters. A naïve method for learning a single Bayesian network with hidden variables is to (1) fix the structure of the Bayesian network (2), use the expectation 30 maximization (EM) algorithm to find good (e.g., ML or

-33-

MAP) parameter values for the Bayesian network and (3) use the parameters obtained from step 2 to compute a score for the model using the Cheeseman-Stutz, BIC or other approximation of the posterior probability for the model. There are two difficulties with this approach. First, the EM algorithm is an iterative algorithm that is too computationally expensive to run on many models. Second, the approximate scores for models with hidden variables (C and H in the case of a mixture of Bayesian networks) do not factor into scores for individual nodes. If it did factor one could use previously calculated scores to make search more efficient. Both problems are solved in the present invention by interleaving the EM algorithm's search for parameters with a search for the structure of the Bayesian networks. By interleaving the search for Bayesian networks and the search for parameters, the invention (in essence) creates scores that factor according to the model and thus allows for efficient search of model structure. In addition, the method of the invention independently searches for each of the Bayesian networks in the mixture of Bayesian networks.

Let  $H_c$  and  $O_c$  be continuous variables (denoted by  $\Gamma_1$  to  $\Gamma_{nc}$  and use  $\gamma_1$  to  $\gamma_{nc}$  to denote values for these variables) and let  $C$ ,  $H_d$  and  $O_d$  be discrete variables (denoted by  $\Delta_1$  to  $\Delta_{nd}$  and use  $\delta_1$  to  $\delta_{nd}$  to denote values for these variables) where  $nc$  is the number of continuous variables and  $nd$  is the number of discrete variables. Let  $\Gamma$  denote the set of

-34-

all of the continuous variables and  $\Delta$  denote the set of all of the discrete variables. We use  $\gamma$  to denote a vector of values for the variables in  $\Gamma$  and  $\delta$  is an index to a configuration of the discrete variables  $\Delta$ .

5 Let  $y_{\text{case}}$  be the configuration of the observed variables  $O$  in a particular case, and let  $x_{\text{case}}$  be a complete configuration of  $C$ ,  $H$  and  $O$ . The key idea to our solution is a concept called complete model sufficient statistics. The complete model sufficient  
10 statistics for a complete case is a vector  $T(x_{\text{case}})$ . This vector is defined as follows:

$$T(x_{\text{case}}) = \langle \langle N_1, R_1, S_1 \rangle, \dots, \langle N_m, R_m, S_m \rangle \rangle$$

From the foregoing definition, the vector  $T(x_{\text{case}})$  consists of  $m$  triples, where  $m$  is the number of  
15 possible discrete configurations for the discrete variables  $\Delta$ . Suppose the discrete variables in  $x_{\text{case}}$  takes on the  $i^{\text{th}}$  configuration. The entries  $N_j$  ( $j \neq i$ ) are zero. The  $R_j$  are vectors of length  $nc$  and the  $S_j$  are square matrices of size  $nc \times nc$ . The  $R_j = \mathbf{0}$  if  
20  $j \neq i$  and  $R_i = \gamma$  otherwise. The  $S_j = \mathbf{0}$  if  $j \neq i$  and  $S_i = \gamma' * \gamma$  otherwise (where  $\gamma'$  is the transpose of  $\gamma$ ). (Note that a boldface zero, e.g.,  $\mathbf{0}$  denotes either a zero vector or matrix.)

-35-

Example involving complete data:

The following is a working example in which a complete database with 2 cases is given as:

5

O1	O2	O3	H1	C
5.1	10	0	1	1
2.7	9	0	0	0

The variables O1 and O2 are continuous. The remaining variables are discrete.

10

In the invention, all possible configurations of the discrete variables are indexed in some fixed way, an example of which is given in the table below:

$\Delta$	C	H1	O3
1	0	0	0
2	0	0	1
3	0	1	0
4	0	1	1
5	1	0	0
6	1	0	1
7	1	1	0
8	1	1	1

-36-

From the foregoing tables, the complete model statistics vector for case 1 is:

$$T(case1) = \left\langle \langle 0,0,0 \rangle, \langle 0,0,0 \rangle, \langle 0,0,0 \rangle, \langle 0,0,0 \rangle, \langle 0,0,0 \rangle, \langle 0,0,0 \rangle, \left\langle 1, \begin{bmatrix} 5.1 \\ 10 \end{bmatrix}, \begin{bmatrix} 26.01 & 51 \\ 51 & 100 \end{bmatrix} \right\rangle, \langle 0,0,0 \rangle \right\rangle$$

5                   The complete model statistics vector for case 2 is:

$$T(case2) = \left\langle \left\langle 1, \begin{bmatrix} 2.7 \\ 9 \end{bmatrix}, \begin{bmatrix} 7.29 & 24.3 \\ 24.3 & 81 \end{bmatrix} \right\rangle, \langle 0,0,0 \rangle, \langle 0,0,0 \rangle, \langle 0,0,0 \rangle, \langle 0,0,0 \rangle, \langle 0,0,0 \rangle, \langle 0,0,0 \rangle, \langle 0,0,0 \rangle \right\rangle$$

The expected complete model sufficient statistics is a vector ECMSS where

10

$$ECMSS = \sum_{case=1}^I E(T(x_{case}) | y_{case}, \theta, S)$$

15                   The expectation of  $T(x_{case})$  is computed by performing inference in a Bayesian network using conventional techniques well-known to those skilled in the art. The sum of  $T(x_1)$  and  $T(x_2)$  is simply scalar, vector, or matrix addition (as appropriate) in each coordinate of the vector.

#### Example involving incomplete data:

20

The following is a working example in which an incomplete data base is given. The incomplete data base is given in the following table, in which the variables O1 and O2 are continuous and O3, H1, and C

-37-

are binary discrete and the symbol "?" denotes unknown data:

O1	O2	O3	H1	C
5.1	10	0	?	?
2.7	9	0	?	?

The vectors  $T(x_{case})$  for each case are readily inferred from the foregoing table in accordance with the definition of  $T(x_{case})$  as follows:

$$T(case1) = \left\langle \left\langle P(\Delta=1|y_1), P(\Delta=1|y_1) \begin{bmatrix} 5.1 \\ 10 \end{bmatrix}, P(\Delta=1|y_1) \begin{bmatrix} 26.01 & 51 \\ 51 & 100 \end{bmatrix} \right\rangle, \dots, \right. \\ \left. \left\langle P(\Delta=m|y_1), P(\Delta=m|y_1) \begin{bmatrix} 5.1 \\ 10 \end{bmatrix}, P(\Delta=m|y_1) \begin{bmatrix} 26.01 & 51 \\ 51 & 100 \end{bmatrix} \right\rangle \right\rangle$$

$$T(case2) = \left\langle \left\langle P(\Delta=1|y_2), P(\Delta=1|y_2) \begin{bmatrix} 2.7 \\ 9 \end{bmatrix}, P(\Delta=1|y_2) \begin{bmatrix} 7.29 & 24.3 \\ 24.3 & 81 \end{bmatrix} \right\rangle, \dots, \right. \\ \left. \left\langle P(\Delta=m|y_2), P(\Delta=m|y_2) \begin{bmatrix} 2.7 \\ 9 \end{bmatrix}, P(\Delta=m|y_2) \begin{bmatrix} 7.29 & 24.3 \\ 24.3 & 81 \end{bmatrix} \right\rangle \right\rangle$$

Having the expected complete model sufficient statistics, the invention uses these as complete model sufficient statistics to perform search among alternative Bayesian networks using the methods described below. The way to do this is to form the expected complete model sufficient statistics for each value of C.

Hence, for each value of C, the expected complete model sufficient statistics for O and H is formed, which is denoted ECMSS\_c. The expected complete model sufficient statistics for O and H can then be used for searching for Bayesian networks. Since the expected complete model sufficient statistics for each

-38-

value of C are distinct (and we have assumed parameter independence) we can use the statistics for each value of C to search for the respective Bayesian network independently of other Bayesian networks. By creating the complete model sufficient statistics we have (in essence) created new scores that factor according to the Bayesian networks, as discussed below in this specification.

For instance, let the indexing of the discrete configurations be as described in the table below.

$\Delta$	C	H1	O3
1	0	0	0
2	0	0	1
3	0	1	0
4	0	1	1
5	1	0	0
6	1	0	1
7	1	1	0
8	1	1	1

Using the index  $\Delta$ , ECMSS<sub>j</sub> is derived from ECMSS by selecting the appropriate triples from ECMSS. For this index we would have

ECMSS<sub>0</sub> = < triple<sub>1</sub>, triple<sub>2</sub>, triple<sub>3</sub>, triple<sub>4</sub> >

ECMSS<sub>1</sub> = < triple<sub>5</sub>, triple<sub>6</sub>, triple<sub>7</sub>, triple<sub>8</sub> >



-39-

In this case, the triple\_j is the triple  $\langle \square_j, R_j, S_j \rangle$  from ECMSS. Specifically, for example, triple\_1 is:

$$\begin{aligned}
 & \left\langle P(\Delta=1|y_2) + P(\Delta=1|y_1)P(\Delta=1|y_2) \begin{bmatrix} 2.7 \\ 9 \end{bmatrix} \right. \\
 & \quad \left. + P(\Delta=1|y_1) \begin{bmatrix} 5.1 \\ 10 \end{bmatrix}, P(\Delta=1|y_2) \begin{bmatrix} 7.29 & 24.3 \\ 24.3 & 81 \end{bmatrix} + P(\Delta=1|y_1) \begin{bmatrix} 26.01 & 51 \\ 51 & 100 \end{bmatrix} \right\rangle \\
 & = \left\langle P(\Delta=1|y_2) + P(\Delta=1|y_1) \begin{bmatrix} 2.7 \times P(\Delta=1|y_2) + 5.1 \times P(\Delta=1|y_1) \\ 9 \times P(\Delta=1|y_2) + 10 \times P(\Delta=1|y_1) \end{bmatrix}, \right. \\
 & \quad \left. \begin{bmatrix} 7.29 \times P(\Delta=1|y_2) + 26.01 \times P(\Delta=1|y_1) & 24.3 \times P(\Delta=1|y_2) + 51 \times P(\Delta=1|y_1) \\ 24.3 \times P(\Delta=1|y_2) + 51 \times P(\Delta=1|y_1) & 81 \times P(\Delta=1|y_2) + 100 \times P(\Delta=1|y_1) \end{bmatrix} \right\rangle
 \end{aligned}$$

From the foregoing, a general process for learning a mixture of Bayesian networks (MBN) in accordance with a exemplary embodiment of the invention is as follows:

1. Choose the number of possible states for the variables C and Hd.
2. Initialize a hypothesis-specific Bayesian-network structure for each hypothetical value of C to be a graph in which each variable in H points to each variable in O --except for the restriction that no continuous variable may point to a discrete variable-- and in which there are no additional arcs. Choose initial values for the parameters in each of these hypothesis-specific Bayesian networks. The parameter values can be set at random, with agglomerative methods, marginal+noise, or other methods. Choose values for the parameters  $\theta_c$ , e.g., choose them to be

-40-

uniform. The marginal+noise initialization method is as follows: Take the initial MBNs, which is a Bayesian network with discrete hidden variables. Remove all hidden nodes and adjacent arcs and  
5 adjust the local distributions, creating model  $s_i$  (a submodel induced by the non-hidden variables). Data is complete with respect to  $s_i$ . Compute MAP parameter values for  $s_i$  (or a model that encodes more independencies than  $s_i$ ). Those practiced in  
10 the art will recognize that this step can be performed in closed form assuming conjugate priors are used. Create a conjugate distribution for the parameters of  $s_i$ ,  $\theta_i$ , whose MAP parameters agree with the MAP parameters just computed and whose  
15 equivalent sample size is specified by the user. This sample size may be different than the one(s) used to determine the parameter priors. Next, for each non-hidden node  $X$  in  $s$  and for each configuration of  $X$ 's hidden parents, initialize  
20 the parameters of the local distributions  $p(x|\Pi_x, \theta_s, s)$  by drawing from the distribution for  $\theta_i$  just described. For each hidden node  $H$  in  $s$  and for each configuration of  $H$ 's (possible) parents, initialize  $H$ 's multinomial parameters to  
25 be some fixed distribution (e.g., uniform). In an alternative embodiment, initialize by sampling from a Dirichlet distribution specified by the user. Those practiced in the art will recognize that this initialization method can be applied to  
30 any parameter optimization algorithm that requires

-41-

- an initial seed (e.g, MCMC, simulated annealing, EM, gradient ascent, conjugate gradient, Newton--Raphson, and quasi-Newton).
3. Use the EM algorithm to do one E step and one M  
5 step to improve the parameter estimates for the current model.
  4. If some convergence criterion is not satisfied then go to step 3.
  5. Using the current MBN, create the expected  
10 complete-model sufficient statistics ECMSS and ECMSS\_c for each hypothesis-specific Bayes net corresponding to  $C=c$ . For every  $C=c$ , translate ECMSS\_c to expected sufficient statistics  $N_{ijk}$ , sample mean, scatter matrix, and sample size (for  
15 use in the structure search step that follows). Those practiced in the art will recognize that this step can be performed with standard techniques.
  6. Using the expected complete model sufficient  
20 statistics for each value of  $C$ , search for structures that improve the score. The result is a new network structure  $s$  with new parameter values.
  7. If some Convergence criterion is not satisfied then go to step 3.
  - 25 8. Save the model that is selected. Choose another number of possible states for the variable  $C$  and  $H_d$ . Go to step 2. Repeat this step and compare the models that are selected. Use the corrected version of the Cheeseman-Stutz score.

-42-

9. The choice in step 4 of whether to go to step 3 or  
step 5 can be decided in a variety of ways, e.g.,  
by checking the convergence of the likelihood, by  
limiting the number of iterations allowed, and so  
5 forth. There are various modifications of this  
process that we have found to be useful including  
having the process adaptively prune out  
hypothesis-specific Bayesian networks for which  $\theta_c$   
(the support for the HSBN corresponding to  $C=c$ )  
10 falls below some threshold (e.g.,  $1/N$ ).

The following is a description of different  
modes of the foregoing process.

15 ((EM)<sup>#</sup>ES\*M)\*:

In this mode of the process, the EM step is  
iteratively repeated (steps 3 and 4) a limited number  
(#) of times while the remainder of the process  
20 including the search for optimum structure is carried  
out to convergence.

((EM)<sup>\*</sup>ES\*M)\*:

25 In this mode of the process, the EM steps are  
iterated until convergence before performing the  
remainder of the algorithm and the structure search is  
also carried to convergence.

-43-

((EM) <sup>#(iteration)</sup> ES\*M) \*:

In this version of the process, the iteration of the EM step (steps 3 and 4) is carried out over a limited number (#) of iterations which is a function of the number of iterations of the structure search step (step 6).

((EM) <sup>#</sup> ES#M) \*:

10

In this version of the process, the number of iterations of the EM step is a fixed number, the number of iterations of the structure search step is a fixed (possibly different) number.

15

((EM) <sup>\*</sup> ES#M) \*:

In this version of the process, the EM steps are always iterated to convergence, while the structure search is iterated a limited number of times.

20

((EM) <sup>#(iteration)</sup> ES#M) \*:

In this version of the process, the number of iterations of the EM step is a function of the number of iterations of the structure search step performed thus far, while the number of iterations of the structure search is a fixed number.

25

The foregoing example uses discrete variables in the Bayesian network where all of the conditional

30

-44-

probability in the Bayesian network are represented as full tables. In an embodiment of the invention described below in this specification, there are decision graphs instead of the tables.

5

#### Implementations of Mixtures of Bayesian Networks

Figure 10 illustrates a Bayesian network consisting of the class variable  $C$  connected to every other variable, the continuous hidden variables  $H_c$  connected to all continuous observed variables, the discrete hidden variables  $H_d$  connected to all observed variables, the continuous observed variables  $O_c$  and the discrete observed variables  $O_d$ . The present invention represents the model depicted in Figure 10 as a mixture of individual Bayesian networks, each individual network corresponding to the hypothesis that the class variable  $C$  is in a particular one of its states (i.e.,  $C=c_i$ ). Each individual network in the mixture is therefore referred to as a hypothesis-specific Bayesian network (HSBN). The corresponding mixture of Bayesian networks (MBNs) consisting of plural HSBNs is illustrated in Figure 11. As indicated in the drawing of Figure 11, one HSBN corresponds to the hypothesis that  $C=c_i$  while another HSBN corresponds to the hypothesis that  $C=c_{i+1}$ , and so forth. In each HSBN of Figure 11, the class variable  $C$  is not included because its state is hypothetically known, and is therefore not a variable. The other variables of the network of Figure 10, namely the hidden variables  $H$  and the

30

-45-

observed variables  $O$  are included in each HSBN of Figure 8. However, after the individual HSBN structures and parameters have been learned, different HSBNs will tend to have different structures, as indicated in Figure 11.

Figure 12 illustrates a first exemplary embodiment of the process for generating mixtures of Bayesian networks (MBNs) discussed above. The first step (block 22 of Figure 12) is to choose the number of states of the external class variable  $C$  and of each discrete hidden variables  $H_d$ . The number of states of  $C$  determines the number of HSBNs in the MBN to be generated. Preferably, when this step is initially performed, the number of states of  $C$  and the number of states of each discrete hidden variable  $H$  are set to their smallest values. For example, if the possible number of states of  $C$  lies in a range of 5 and 10, the number of states of  $H_{d1}$  lies in a range of 3 and 6 and the number of states of  $H_{d2}$  lies in a range of 11 and 14, then the lowest number in each range is chosen initially. In subsequent repetitions of this step by an outer loop (which will be discussed below), all combinations of the numbers of states are eventually chosen.

The next step (block 24 of Figure 12) is to initialize an MBN. Preferably, this is done by forming an MBN consisting of identical HSBNs with an arc from each hidden variable to each observed variable with the

-46-

proviso that continuous hidden are connected only to continuous observed variables as shown in Figure 11. Also, in this step the HSBN parameters are initialized using the marginal+noise method. The

5 expectation-maximization step (block 26 of Figure 12) is then performed on all HSBNs in the MBN. The expectation-maximization step is described in Dempster et al, "Maximum Likelihood From Incomplete Data Via the EM Algorithm", Journal of the Royal Statistical

10 Society B, Volume 39 (1977). This step produces a more optimal version of the parameters of the individual HSBNs. A test for convergence is then performed (block 28 of Figure 12). If the

15 expectation-maximization step has not converged (NO branch of block 28), then the process loops back to the expectation-maximization step of block 26 in an inner loop (loop 2 of Figure 12). Otherwise (YES branch of block 28), the network parameters are saved (block 30).

20 The expected complete-model sufficient statistics (ECMSS) are then computed (block 32). The computation of each of the probabilities  $p(\Delta)$  in  $T(x_{\text{case}})$  is performed by conventional inference techniques using the current version of the MBN. How inferencing is

25 performed with an MBN is described below herein with reference to Figure 16. The computation of  $T(x_{\text{case}})$  has been described above in this specification. The ECMSS are then translated (block 34 of Figure 12) using conventional techniques into expected sufficient

30 statistics  $N_{ijk}$ , sample means, scatter matrix, and sample



-47-

size for each HSBN, all of which are defined below in this specification with reference to a structure search process. Next, an optimum structure is found for each HSBN (block 36) by treating the expected sufficient

5 statistics as sufficient statistics for complete data. The step of block 36 includes searching for the optimum structure of the HSBN (block 38) and saving the optimum structure HSBN (block 40). The search of block 38 is described below in this specification, and employs the

10 expected sufficient statistics, sample means, scatter matrix and sample size computed from the ECMSS in the step of block 34. The search is based upon scoring each candidate structure of the HSBN, the score being the marginal likelihood of the expected complete data  $D$

15 given the candidate network structure  $s$ , namely  $p(D|s)$ . With each selection of optimal structures for the HSBNs of the MBNs, an overall MBN score is as follows (block 42):

$$20 \quad \text{score}(s) = p(s)p(D_c | s) \frac{p(D | \tilde{\theta}, s)}{p(D_c | \tilde{\theta}, s)} \quad (1)$$

where  $\tilde{\theta}$  denotes the MAP parameters given  $D$ ,  $D_c$  is a complete data set whose sufficient statistics are equal to the expected complete-model sufficient statistics,

25 and  $p(s)$  is the prior probability of MBN structure  $s$  (prior on structure). The prior on structure  $p(s)$  is given by:

$$p(s) = p(|C|) \prod_{B'_s \in MBN} p(B'_s | |C|)$$

-48-

where  $|C|$  is the number of states of the hidden variable  $C$ . The exemplary embodiment uses the log of the score in order to reduce numerical instability:

$$\begin{aligned} \log score(s) = & \log p(|C|) + \sum_{B'_i \in MBN} \log p(B'_i || C) + \log p(D_c | s) \\ & + \log p(D | \tilde{\theta}, s) - \log p(D_c | \tilde{\theta}, s) \end{aligned}$$

This MBN score is the Cheeseman-Stutz score. (See P. Cheeseman and J. Stutz, "Bayesian Classification AutoClass: Theory and Results", Advances in Knowledge Discovery and Data Mining, AAAI Press [1995]).

Next, a test for convergence of the structure search of step 38 is performed (block 44). The test for convergence in this embodiment consists of inquiring whether any HSBN structure within the present MBN has changed since the performance of this convergence test or since the HSBNs were initialized. If there have been any structural changes (YES branch of block 44), the structure search has not converged and the process loops back to the expectation-maximization step of block 26 in loop 1 of Figure 12. Otherwise (NO branch of block 44), with no structural changes since the previous iteration of loop 1, the structure search has converged and the next step is to determine whether the various combinations of the number of states of the discrete class variable and discrete hidden variables has been exhausted (block 46 of Figure 12). If not (NO branch of block 46), the process loops back in an outer loop (loop 0 of Figure 12) to the step of block 22 in which

-49-

the next combination of number of states is selected. Otherwise (YES branch of block 46), the MBN having the highest score is selected and output for use in performing inferencing (block 48). Alternatively, some  
5 or all of the MBNs are output as a collection of MBNs along with their respective MBN scores. In this alternative mode (described below herein), inferencing from a given input is performed by all of the MBNs in the collection in parallel, their outputs being weighted  
10 in accordance with their respective MBN scores, and the weighted sum of the MBN outputs being the output of the collection of MBNs.

Figure 13 illustrates an alternative  
15 embodiment in which the test for convergence of block 44 consists of determining whether the MBN score has decreased since the previous iteration of the middle loop. If the MBN score has not increased, then loop 1 of Figure 13 has converged. Otherwise, if the score has  
20 increased then loop 1 has not converged.

Figure 14 illustrates a variation of the embodiment of Figure 12 in which the number of iterations  $T$  of the inner loop is a function  $T(S)$  of the  
25 number of iterations  $S$  of the outer loop. The first step (block 22 of Figure 14) is to choose the number of states of the external class variable  $C$  and of each discrete hidden variables  $H_d$ . The next step (block 24 of Figure 14) is to initialize an MBN. Then, the number  
30 of iterations of the outer loop,  $S$ , is initialized to

-50-

zero (block 50). The expectation-maximization step (block 26 of Figure 14) is then performed on all HSBNs in the MBN. A determination is then made (block 28' of Figure 14) of whether the expectation-maximization process has converged, or, if not, whether loop 2 of Figure 14 (the "inner loop") has iterated T(S) times. If neither condition holds (NO branch of block 28'), then the process loops back to the expectation-maximization step of block 26 in the inner loop. Otherwise (YES branch of block 28), a flag is set if the expectation-maximization process has not converged after T(S) iterations of the inner loop (block 52). The network parameters are then saved (block 30).

15

The expected complete model sufficient statistics (ECMSS) are then computed (block 32). The ECMSS are then translated (block 34 of Figure 14) using conventional techniques into expected sufficient statistics  $N_{ijk}$ , sample means, scatter matrix and sample size, all of which are defined below in this specification with reference to a structure search process. Next, an optimum structure is found for each HSBN (block 36). The step of block 36 includes searching for the optimum structure of the HSBN (block 38) and saving the optimum structure HSBN (block 40). The search of block 38 is described below in this specification, and employs the expected sufficient statistics, sample means, scatter matrix and sample size computed from the ECMSS in the step of

30

-51-

block 34. The search is based upon scoring each candidate structure of the HSBN, the score being the marginal likelihood of the expected complete data  $D_c$  given the candidate network structure  $B_i^e$ , namely

5  $P(D_c | B_i^e)$ . With each selection of optimal structures for the HSBNs of the MBN, an overall MBN score is computed as described with reference to Figure 12 (block 42). This MBN score is the corrected version of the MBN score in accordance with the Cheeseman-Stutz score (see P.  
10 Cheeseman and J. Stutz, "Bayesian Classification AutoClass: Theory and Results", Advances in Knowledge Discovery and Data Mining, AAAI Press [1995]).

Next, a test for convergence of the structure  
15 search of step 38 is performed (block 44). The test for convergence in this embodiment consists of inquiring whether any HSBN structure within the present MBN has changed since the performance of this convergence test. If there have been any structural changes (YES branch of  
20 block 44), the structure search has not converged, in which case  $S$  is incremented (block 54) while the process loops back to the expectation-maximization step of block 26 through loop 1 of Figure 14. Otherwise (NO branch of block 44), with no structural changes since  
25 the previous iteration of loop 1 or 1', the structure search has converged and the next step is to determine whether the flag is set (block 56). If so (YES branch of block 56), the flag is reset (block 58),  $S$  is incremented (block 54') and the process loops back to  
30 the expectation maximization step of block 26 through

-52-

loop 1' of Figure 14. Otherwise, if the flag is not currently set (NO branch of block 56), a determination is made of whether the various combinations of the number of states of the discrete class variable and discrete hidden variables has been exhausted (block 46 of Figure 14). If not (NO branch of block 46), the process loops back in the outer loop (loop 0 of Figure 14) to the step of block 22 in which the next combination of number of states is selected. Otherwise (YES branch of block 46), the MBN having the highest score is selected and output for use in performing inferencing (block 48).

Figure 15 illustrates a modification of the embodiment of Figure 14 in which the test for convergence of block 44 consists of determining whether the MBN score has decreased since the previous iteration of the middle loop. If the MBN score has decreased, the loop 1 of Figure 12 has converged. Otherwise, if the score has not decreased loop 1 has not converged.

Figure 16 illustrates an inferencing apparatus including an MBN. The MBN includes a set of HSBNs 60, each of which is associated with a weight 62 equal to the probability of the class variable C being in the corresponding state. Multipliers 64 combine the output of each HSBN with the corresponding weight 62 and an adder 66 computes the sum of the products. An input is applied to all the HSBNs 60 simultaneously, resulting in a single inference output from the MBN.

-53-

Figure 17 illustrates an inferencing apparatus including a collection of MBN's. Each MBN is of the type described in Figure 16. Of course, the scores have been previously computed as described above at the time each MBN is generated before inferencing is performed. Each MBN output is weighted by the corresponding MBN score by a multiplier 72, and an adder 74 combines the weighted MBN outputs into a single output of the collection of MBNs.

How to Perform the Structure Search Step of Block 38 --  
Searching for Optimum Structure and Scoring the  
Hypothesis-Specific Network Structure:

Figure 18 depicts a diagram of the MBN generator 502 of the exemplary embodiment of Figure 6. The MBN generator 502 of the exemplary embodiment contains a scoring mechanism 602 and a network adjuster 606. The scoring mechanism 602 receives the expert knowledge 506, the empirical data 504, the test network 608 and a list of nodes 610 as input. After receiving this information, the scoring mechanism 608 generates a score 604 that ranks the nodes of test network 608 as indicated by the list of nodes 610 for goodness. Thus, the score 604 contains a subscore for each node scored. Each subscore indicates how well the portion of the test network involving the node corresponding to the subscore and the parents of the node is at rendering inferences based on the empirical data 504 and the expert knowledge 506. The test network 608 received as input is either the prior

-54-

network or a test network 608 generated by the network adjuster 606 depending on the circumstances. That is, the scoring mechanism 602 of the exemplary embodiment uses the initial network as the test network for the first invocation of the scoring mechanism. After the first invocation of the scoring mechanism 602, the test network received by the scoring mechanism is the test network 608 generated by the network adjuster. In the exemplary embodiment, a Bayesian network (i.e., the initial network or the test network 608) is stored in memory as a tree data structure where each node in the tree data structure corresponds to a node in the Bayesian network. The arcs of the Bayesian network are implemented as pointers from one node in the tree data structure to another node. In addition, the probabilities for each node in the Bayesian network are stored in the corresponding node in the tree data structure.

The network adjuster 606 receives as input the score 604 and the initial network and generates a new test network 608 in response thereto, which is then passed back to the scoring mechanism 602 with a list of nodes 610 which need to be rescored. After iterating many times between the scoring mechanism 602 and the network adjuster 606, the network adjuster eventually generates an improved MBN 508 (hereinafter referred to as a Bayesian network). The network adjuster 606 generates the improved Bayesian network 508 when the scores 604 generated do not improve. That is, the



-55-

network adjuster 606 retains the test network 608 that the network adjuster last generated, modifies the test network based on the score 604, and if the network adjuster cannot generate a test network with a better score than the retained test network, the network adjuster generates the retained test network as the improved Bayesian network 508. Although the exemplary embodiment has been described as iterating many times between the scoring mechanism 602 and the network adjuster 606, one skilled in the art will appreciate that only one iteration may be performed. The initial network used by the scoring mechanism 602 of the exemplary embodiment can consist of all discrete variables, all continuous variables, or a combination of discrete and continuous variables.

Figure 19 depicts a high level flow chart of the steps performed by the scoring mechanism 602 of the exemplary embodiment. The scoring mechanism 602 of the exemplary embodiment determines the types of variables used in the test network 608 and generates a score for the test network. First, the scoring mechanism of the exemplary embodiment determines if the test network 608 contains all discrete variables (step 702). If the test network 608 contains all discrete variables, the scoring mechanism 602 generates a score for the nodes in the list of nodes 610 of the test network by invoking the calculate discrete score process (step 704). However, if the test network 608 does not contain all discrete variables, the scoring

-56-

mechanism 602 determines if the test network contains all continuous variables (step 706). If the test network 608 contains all continuous variables, the scoring mechanism 602 generates a score for the nodes indicated in the list of nodes 610 of the test network by invoking the calculate continuous score process (step 708). However, if the test network 608 does not contain all continuous variables, the test network contains a combination of discrete and continuous variables ("a mixed network"), and the scoring mechanism generates a score for the nodes indicated by the list of nodes 610 of the test network by invoking the calculate mixed score process (step 710).

The calculate discrete score process, the calculate continuous score process and the calculate mixed score process are based upon a common concept, Bayes' theorem. The score that each scoring process produces is proportional to the posterior probability of the test network. That is, probability distributions and densities can be of two types: prior and posterior. The prior probability distribution or density is the probability distribution or density before data is observed. The posterior probability distribution or density is the probability distribution or density after data is observed. Bayes' theorem states that the posterior probability of a test network is proportional to the prior probability of a test network multiplied by the probability of the empirical

-57-

data database given the test network and the expert knowledge.

Calculate Discrete Score Process:

5

The calculate discrete score process scores a test network containing all discrete variables. The calculate discrete score process takes advantage of the fact that the probability of the empirical data database given a test network and expert knowledge is the product over all cases of the probability of a particular case given a test network, expert knowledge, and previous cases (i.e., cases observed prior to the particular case). The computation of the probability of a case given a test network, expert knowledge, and previous cases is based on the assumption that the empirical data database represents a multinomial sample from the test network. That is, the empirical data database contains a sequence of observations that form a multinomial distribution as described in DeGroot, Optimal Statistical Decisions, at 48-49 (1970). Thus, each variable given each instance of the parents of the variable is associated with a set of parameters  $(\theta_{ij1}, \dots, \theta_{ijr_i})$ , where:

25

$i$  is the variable index, " $i=1 \dots n$ ," where " $n$ " is the number of variables in the test network;

$j$  is the parent-instance index; " $j=1 \dots q_i$ ," where  $q_i$  is the number of instances of the parents;

-58-

k is the variable state index; "k=1...r<sub>i</sub>", where "r<sub>i</sub>" is the number of states of the variable i.

The parameter  $\theta_{ijk}$  is the long run fraction  
 5 for  $x_i = k$ , when  $\Pi_i = j$ . That is, for all values of i, j, and k,  $p(x_i = k | \Pi_i = j, \theta_{ijk}, B_s^e, \xi) = \theta_{ijk}$ , where  $B_s^e$  is the hypothesis-specific test network.

In addition, the exemplary embodiment assumes  
 10 that the density of each parameter set  $\{\theta_{ij1}, \dots, \theta_{ijr_i}\}$  has a Dirichlet distribution as defined by:

$$\rho(\theta_{ij1}, \dots, \theta_{ijr_i} | B_s^e) = \frac{\Gamma\left(\sum_{k=1}^{r_i} (e_i)\right)}{\prod_{k=1}^{r_i} \Gamma(e_i)} \prod_{k=1}^{r_i} \theta_{ijk}^{e_i-1}, e_i > 0$$

15 where " $\Gamma()$ " is the Gamma function defined as

$$\Gamma(x) = \int_0^{\infty} e^{-y} y^x dy. \quad \text{The exponents } e_i \text{ are given by } K/(r_i q_i)$$

where K is the sample size specified by the user.  
 Alternatively, one may use  $e_i = 1$ .

20 Figure 20 depicts a flow chart of the steps performed by the calculate discrete score process. The first step of the calculate discrete score process is to examine the translated expected complete-model sufficient statistics ECMSS\_c for the number of times  
 25 ("hits") that each variable is encountered, for each

-59-

state of each variable, and for each possible instance of the parents of each variable. The number of hits, therefore, has three indices  $i$ ,  $j$  and  $k$ ; " $i=1...n$ ," where " $n$ " is the number of variables in the test network; " $j=1...q_i$ ," where  $q_i$  is the number of instances of the parents; and " $k=1...r_i$ ," where " $r_i$ " is the number of states of variable  $i$ . Next, the calculate discrete score process of the exemplary embodiment selects a variable from the test network 608 according to the list of nodes 610 to score, starting with the first variable in the list of nodes (step 806). After a variable is selected, the calculate discrete score process calculates a subscore for the selected variable (step 808) and stores the calculated subscore in the node of the test network that corresponds with the selected variable (step 810).

The subscore for each variable  $x_i$  is calculated using the following formula:

$$\log \left[ \prod_{j=1}^{q_i} \frac{\Gamma(r_i e_i)}{\Gamma(N_{ij} + r_i e_i)} \prod_{k=1}^{r_i} \frac{\Gamma(N_{ijk} + e_i)}{\Gamma(e_i)} \right]$$

After the subscore for a variable is calculated and stored, the calculate discrete score process determines in step 812 if there are more variables to be processed and either continues to step 806 to process more variables, or continues to step. After storing the subscores, a total score for the test network is generated by adding all of the subscores

-60-

together and adding the log prior probability of the HSN structure given  $|C|$  (step 822).

Calculate Continuous Score Process:

5

The calculate continuous score process calculates scores for test networks containing all continuous variables and is based on Bayes' theorem. The calculate continuous score process assumes that all cases in the empirical data database are drawn from a multivariate normal distribution. The calculate continuous score process takes advantage of the fact that a set of variables have a multivariate normal distribution if and only if each particular variable is an independent (univariate) normal distribution, when conditioned on the variables that precede the particular variable in some ordering:

10

15

$$p(x_i | x_1 \dots x_{i-1}) = n \left( m_i + \sum_{j=1}^{i-1} b_{ji} (x_j - m_j), 1/v_i \right)$$

20

The term  $p(x_i | x_1 \dots x_{i-1})$  denotes the density of a particular variable given all the variables before the particular variable in some ordering. The term

25

$$n \left( m_i + \sum_{j=1}^{i-1} b_{ji} (x_j - m_j), 1/v_i \right)$$

contains "n" referring to a normal distribution having a mean " $m_i$ ", a variance " $v_i$ " and coefficients " $b_{ji}$ ".

-61-

"m", "v" and "b" are parameters of the normal distribution. The coefficient " $b_{ji}$ " refers to the strength of the connection between the mean of a variable " $X_i$ " and the value of the variables " $X_j$ ".

5 Thus,  $b_{ji}$  is equal to zero if and only if there is no arc from " $X_j$ " to " $X_i$ " in the test network. One skilled in the art would recognize that the coefficient " $b_{ji}$ " is sometimes called a partial regression coefficient between " $x_i$ " and " $x_j$ ". The multivariate normal  
10 distribution and the univariate normal distribution are well known in the field of statistics.

In addition, the calculate continuous score process is based on three assumptions. First, the  
15 calculate continuous score process assumes that the prior distributions for the mean and precision matrix of the multivariate normal distribution with all dependencies between variables being possible (i.e.,  $B_{sc}^e$ ) is the normal-Wishart distribution. The  
20 normal-Wishart distribution is described in DeGroot, Optimal Statistical Decisions, at 56-59 (1970). The normal-Wishart distribution is conjugate for multivariate normal sampling. Second, the parameters  $E((v_1b_1), \dots, (v_nb_n))$  are mutually  
25 independent. Third, if  $x_i$  has the same parents in two different Bayesian networks, then the prior densities of "v" and "b" of  $X_i$  for both Bayesian networks are the same.

-62-

Figures 21A and 21B depict a flow chart of the steps performed by the calculate continuous score process of the exemplary embodiment. The calculate continuous score process of the exemplary embodiment first calculates the parameters associated with the prior densities of the normal-Wishart distribution as follows (step 902):

$$T_0 = I$$

$$\bar{\mu}_0 = \langle \bar{x} \rangle$$

10

where " $T_0$ " is the precision matrix of the normal-Wishart distribution (an  $n$  by  $n$  matrix),  $I$  is the identity matrix, " $\bar{\mu}_0$ " is the prior mean of the normal-Wishart distribution (an  $n$  by 1 column matrix), and " $\langle \bar{x} \rangle$ " is the sample mean of the variables in the domain. The calculate continuous score process then examines the sufficient statistics. That is, the sample mean and the multivariate internal scatter matrix (step 906). For complex data, the sample mean is defined by:

20

$$\bar{x}_m = \frac{1}{m} \sum_{i=1}^m \bar{x}_i$$

25

where " $\bar{x}_m$ " refers to the sample mean, " $m$ " is the number of complete cases in the database, and " $\bar{x}_i$ "



-63-

refers to a case. The multivariate internal scatter is defined by:

$$S_m = \sum_{i=1}^m (\bar{x}_i - \bar{x}_m)(\bar{x}_i - \bar{x}_m)'$$

5 where " $S_m$ " refers to the multivariate internal scatter matrix, where " $\bar{x}_i$ " refers to a case, and where " $\bar{x}_m$ " refers to the sample mean. The mark ' refers to the transpose in which the matrix is rearranged from being an "n by 1" to being a "1 by n" matrix, and multiplied  
10 together so as to render an "n by n" matrix.

The calculate continuous score process next combines the intermediate statistics obtained from steps 902 and 906 (step 908). In this step,  $T_0^{n \times n}$   
15 (indicating that  $T_0$  is an n by n matrix) is combined with the multivariate internal scatter matrix and a term involving the sample mean and prior mean to create  $T_m^{n \times n}$ . In this step, the following is computed:

$$20 \quad T_m^{n \times n} = T_o^{n \times n} + S_m^{n \times n} + \frac{Km}{K+m} (\bar{\mu}_o - \bar{x}_m)(\bar{\mu}_o - \bar{x}_m)'$$

where "K" is the effective sample size specified by the user, "m" is the number of completed cases in the expected complete-model sufficient statistics ECMSS\_c,  
25 " $T_0$ " is the precision matrix of the prior normal-Wishart distribution, " $\bar{\mu}_o$ " is the prior mean of

-64-

the normal-Wishart distribution, and " $\bar{X}_m$ " is the sample mean.

Next, the calculate continuous score process of the exemplary embodiment selects one variable from the list of nodes to be scored (step 910). After selecting one variable, the calculate continuous score process calculates a subscore ("the complete data subscore") for that variable and stores the complete data subscore into the node (step 912). The calculate continuous score process calculates the subscore for one variable by performing the following:

$$\text{subscore}[i] = \log \left[ p(B_i^e(i) | \xi) \frac{p(D^{x_i \Pi_i} | B_i^e)}{p(D^{\Pi_i} | B_i^e)} \right]$$

15

The term " $p(B_i^e(i) | \xi)$ " refers to the prior probability of the variable-parent pair  $x_i - \Pi_i$ . Both terms in the fraction are computed using

$$p(D^R | B_{sc}^e) = (2\pi)^{-nm/2} \left( \frac{K}{K+m} \right)^{n/2} \frac{c(n, K+n)}{c(n, K+n+m)} |T_0|^{(K+n)/2} |T_m|^{\frac{K+n+m}{2}}$$

20

The term " $p(D^R | B_{sc}^e)$ " refers to the density of the data restricted to the set of variables R given the event indicated by the prior network  $B_{sc}^e$ , where "n" is the number of variables in R, "K" is the effective sample size specified by the user, "m" is the number of completed cases in the ECMSS\_c, " $|T_0|$ " is the

25

-65-

determinant of  $T_0$  marginalized to the variables in  $R$ , " $|T_m|$ " is the determinant of  $T_m$  marginalized to the variables in  $R$ , and  $c(n, K)$  is the Wishart normalization function defined as:

5

$$c(n, K) = \left\{ 2^{Kn/2} \pi^{n(n-1)/4} \prod_{i=1}^n \Gamma\left(\frac{K+1-i}{2}\right) \right\}^{-1}$$

The determinant of an  $n$  by  $n$  matrix ( $A$ ) is the sum over all permutations  $p = (i_1 \dots i_n)$  of the integers 1 through  $n$  of the product:

10

$$(-1)^{k_p} \prod_{j=1}^n A[j, i_j]$$

where  $k_p$  is 0 if  $P$  is even and  $k_p$  is 1 if  $P$  is odd.

15

After the calculate continuous score process of the exemplary embodiment calculates a subscore for one variable, the calculate continuous score process determines if there are more variables to be processed (step 914). If there are more variables in the list of nodes for processing, the calculate continuous score process continues to step 910. However, if there are no more variables for processing in the test network, the calculate continuous score returns. Finally, the calculate continuous score process calculates the total score by adding all the subscores together and adding the log prior probability of the hypothesis-specific network structure given  $|C|$  (step 922).

25

-66-

Calculate Mixed (Discrete and Continuous) Score Process:

5       The calculate mixed score process calculates  
a score for a mixed network having both discrete and  
continuous variables, and is based on Bayes' theorem.  
In calculating a score for a mixed network, the  
exemplary embodiment enforces a restriction that the  
10       initial network be constructed under the assumption  
that all dependencies among variables are possible.  
This restriction is enforced by the knowledge engineer.  
The exemplary embodiment also enforces a restriction  
that the prior network and all mixed test networks  
correspond to a collection of conditional Gaussian  
15       distributions. This restriction is enforced by the  
knowledge engineer and the network adjuster,  
respectively. In the following discussion, the  
symbols  $\Gamma$ ,  $\Delta$ ,  $\Gamma_i$  and  $\Delta_i$  appearing above in this  
specification are employed here, but have a different  
20       meaning. For the domain of all variables in a mixed  
network to be a collection of conditional Gaussian  
distributions, the set of continuous variables " $\Gamma$ " and  
the set of discrete variables " $\Delta$ " must be divisible  
into disjoint sets  $\Gamma_1 \dots \Gamma_\gamma$  such that for each set  $\Gamma_i$   
25       there exists a  $\Delta_i$  subset of  $\Delta$  such that  $\Gamma_i$  is connected  
with respect to continuous variables,  $\Gamma_i$  and  $\Gamma_j$  ( $i \neq j$ )  
is not connected with respect to continuous variables,  
no continuous variable is the parent of a discrete

-67-

variable, and  $\Delta_i$  is a minimal set such that  $\Gamma_i$  and  $\Delta$  are conditionally independent given  $\Delta_i$ .

Figure 22 depicts the flow chart of the steps performed by the calculate mixed score process of the exemplary embodiment. The effect of the calculate mixed score process of the exemplary embodiment is that, first, the discrete variables are scored. Then, for each subset  $\Delta_i$  and for each instance of subset  $\Delta_i$ , the scores for the continuous variables in  $\Gamma_i$  are calculated and added. Lastly, the log prior probability for the HSNB is added to the score.

The first step that the calculate mixed score process of the exemplary embodiment performs is to calculate the subscore for all discrete variables in the list of nodes to be scored (step 1002). The calculate mixed score process performs this by invoking the calculate discrete score process on the test network restricting the nodes scored to only the discrete nodes. The calculate mixed score process then selects a set of continuous variables " $\Gamma_i$ " from the list of nodes to be scored (step 1004). Next, the calculate mixed score process selects a variable within " $\Gamma_i$ " for processing (step 1006). After selecting a variable, the calculate mixed score process calculates a continuous subscore for the selected continuous variable for all instances of the parents of the variable (step 1008). In calculating the continuous

-68-

subscore for a mixed network, since the mixed network is divided into sets of nodes, the definitions for  $K$ ,  $\mu_0$ , and  $T_0$ , as described relative to the calculate continuous score process are redefined as a function of  $i$  and  $j$  (the instance of  $\Delta_i$ ).

$$K_{ij} = K / q_i$$

$$\bar{\mu}_{ij} = \langle \Gamma_i | \Delta_i = j \rangle$$

$$T_{ij} = I$$

10

where " $q_i$ " is the number of parents of  $\Gamma_i$ , and " $\Delta_i$ " are as defined above " $\bar{\mu}_{ij}$ " is redefined as the sample mean of variables " $\Gamma_i$ " given the discrete parents of " $\Gamma_i$ " that equal configuration  $j$ . Alternatively, each effective sample size  $K_{ij}$  may be specified by the user.

15

The calculate mixed score process then determines if there are more variables in the selected set for processing (step 1010). If there are more variables to be processed, processing continues to step 1006. However, if there are no more variables to be processed, processing continues to step 1012 wherein the calculate mixed score process determines if there are more sets of continuous variables to be processed. If there are more sets of continuous variables to be processed, then processing continues to step 1004. However, if there are no more sets of continuous

20

25

-69-

variables to be processed, then the calculate mixed score process continues to step 1014 wherein the calculate mixed score process adds the discrete subscores, the continuous subscores and the log prior on the HSBN structure together. Steps 1004 through 1014 can therefore be described using the following formula:

$$score(B_s^e) = \log p(B_s^e | C) + \log p(D^\Delta | B_s^e) + \sum_{k=1}^r \sum_{j=1}^{q_k} \sum_{i=1}^n \log \frac{p(D^{x_i \Pi_i} | \Delta_k = j, B_s^e)}{p(D^{\Pi_i} | \Delta_k = j, B_s^e)}$$

where " $\log p(B_s^e | C)$ " refers to the log prior on structure  $B_s^e$ , given  $|C|$ , the term " $\log p(D^\Delta | B_s^e)$ " refers to the score for the discrete variables in the test network, and  $q_k$  is the number of configurations of  $\Delta_i$ . In addition, the term,

$$\sum_{k=1}^r \sum_{j=1}^{q_k} \sum_{i=1}^n \log \frac{p(D^{x_i \Pi_i} | \Delta_k = j, B_s^e)}{p(D^{\Pi_i} | \Delta_k = j, B_s^e)}$$

refers to the score for the continuous variables wherein the term " $D^{x_i \Pi_i}$ " refers to the data restricted to variables  $\{X_i\} \cup \Pi_i$ .

#### Network adjuster:

Figures 23A and 23B depict a flow chart of the steps performed by the network adjuster 606 of the exemplary embodiment of the present invention.

-70-

The network adjuster processes the test network stored on the last invocation of the network adjuster (or a newly created initial network) and selects a node within the test network for processing, starting with the first (step 1102). The network adjuster then performs all legal single changes on the selected node (step 1104). That is, the network adjuster in sequence: adds an arc to the selected node from each other node (not already directly connected) as long as the new arc does not introduce a directed cycle, deletes each arc pointing to the selected node, and reverses each arc pointing to the selected node as long as the modified arc does not introduce a directed cycle. In addition, if the test network is a mixed network, the network adjuster ensures that the test network remains conditional Gaussian. The network adjuster next requests the scoring mechanism to generate new subscores for each legal change for the affected nodes (step 1106). The affected nodes are the nodes at either end of an arc change. Since the data has been completed so that there is no missing data, the exemplary embodiment can perform changes on a node-by-node basis because the subscores of each variable obtained for the discrete variable networks, the continuous variable networks, and the mixed networks, are logically independent. In other words, the score is said to be factorable. Therefore, because the score is factorable, if the subscore for the affected nodes improve, it can be ensured that the entire score will improve. The subscores are generated



-71-

using the calculate discrete score process, the calculate continuous score process, or the calculate mixed score process, depending on the type of the test network. The network adjuster then selects the change  
5 that produces the best subscore for the affected nodes (step 1108).

After the best change for the selected nodes has been identified, the network adjuster of the  
10 exemplary embodiment determines whether there are more variables in the test network for processing (step 1110). If there are more variables in the test network for processing, the network adjuster proceeds to step 1102 wherein the next variable in the test  
15 network is selected for processing. After all of the variables have been processed, the network adjuster identifies the single change of the best changes selected from step 1108 that most improves the total score of the test network (step 1111). If there is  
20 such a change, then the network adjuster stores the test network and the subscores for the affected nodes, and then returns to step 1102. If no change exists that improves the total score, then the network adjuster returns the current test network as the  
25 improved Bayesian network 508.

Preferred Calculating Discrete Score Method -  
Employing Decision Graphs in Each Variable:

30 The calculate discrete score process described above is not the presently preferred

-72-

embodiment. The presently preferred embodiment is described in U.S. application Serial No. \_\_\_\_\_ filed \_\_\_\_\_ entitled "Belief Networks with Decision Graphs". This preferred method is now  
5 described in this specification for use in carrying out the present invention.

An exemplary embodiment of the preferred discrete score calculation utilizes a decision graph in  
10 each of the nodes of a Bayesian network to store the probabilities for that node. A decision graph is an undirected graph data structure where each vertex is connected to every other vertex via a path and where each leaf vertex may have more than one path leading  
15 into it, which forms a cycle. An exemplary decision graph 1400 is depicted in Figure 24. This decision graph 1400 is for a node z of a Bayesian network where node z has parents x and y. As can be seen from decision graph 1400, it contains one root vertex and  
20 only three leaf vertices, because one of the leaf vertices contains a probability for two sets of values: where x equals 0 and y equals 1, and where x equals 1 and y equals 0.

25 A decision graph is a much more flexible and efficient data structure for storing probabilities than either a tree or a table, because a decision graph can reflect any equivalence relationship between the probabilities and because leaf vertices having  
30 equivalent probabilities need not be duplicated.

-73-

Additionally, by being able to reflect an equivalency relationship, multiple paths (or combinations of the parent values) can refer to the same probability, which yields a more accurate probability. For example, if  
5 there are 8 possible combinations of the parent vertices' values, if one probability is stored for each combination, and if the Bayesian network was created using a database of 16 cases, the ratio of cases to probabilities is 2 to 1. A case is a collection of  
10 values for the nodes of the Bayesian network (and, consequently, the vertices of the decision graph) that represents real-world decisions made in a field of decision making. In other words, each probability was created using two data points on average. However, if  
15 the number of probabilities stored is reduced such that more than one combination refers to a probability, the ratio of cases to probabilities improves so that the probability becomes more accurate given the data. That is, some of the probabilities are based on an increased  
20 number of data points, which produces more accurate probabilities.

#### Overview of Decision Graphs

25 An exemplary embodiment of the present invention receives an equivalent sample size, which is the equivalent number of times the expert has provided decision-support in the field of expertise (e.g., the number of times that an automobile mechanic has  
30 diagnosed a particular automobile problem).

-74-

Additionally, the exemplary embodiment receives the ECMSS\_c summarizing many real-world cases. After receiving this information, the exemplary embodiment creates initial decision graphs for the nodes of the hypothesis-specific Bayesian network and then adjusts the decision graphs to better reflect the data. During the learning process, the decision graphs are scored to determine goodness at reflecting the data, and a number of candidate decision graphs are generated for each node by making adjustments to the decision graphs contained in each node. These candidate decision graphs are then scored, and the candidate decision graph with the best score (i.e., the score that improves the most) is stored for each node. After storing the decision graph with the best score into each node, the Bayesian network is scored for how well all decision graphs reflect the data, and the Bayesian network is then updated to improve its score. The adjustments to the Bayesian network include adding arcs between the nodes to reflect additional relationships that were identified during the learning process. The learning process continues until the Bayesian network with the best possible score is produced.

Although the hypothesis-specific Bayesian network of an exemplary embodiment can be used in numerous decision-support systems, it is described below with reference to a particular decision-support system for use in predicting whether a user would like to visit a web site on the Internet based on various

-75-

characteristics of the user. Predicting whether a user would like a particular web site is referred to as web site analysis. A Bayesian network suitable for use in performing such web site in accordance with one example is depicted in Figure 25A. Figure 25A shows a Bayesian network 1500 containing various nodes 1502-1508 and arcs connecting the nodes 1510-1518. The age node 1502 represents the age of the user and has a number of states or values including: 0 for ages 0-18, 1 for ages 19-30, 2 for ages 31-40, and 3 for ages greater than 40. The sex node 1504 contains a value indicating the sex, either male or female, of the user. The business node 1506 contains a value (i.e., 0 for no and 1 for yes) indicating whether a particular user visited business-related web sites, and the travel node 1508 contains a value (i.e., 0 for no and 1 for yes) indicating whether a particular user visited travel-related web sites. As can be seen from arcs 1510-1516, the values of both the age node 1502 and the sex node 1504 influence whether the user would like to visit business-related web sites as reflected by node 1506 as well as whether the user would like to visit travel-related web sites as reflected by node 1508. Additionally, the value of the business node 1506 influences the value of travel node 1508. An exemplary embodiment uses the Bayesian network 1500 to perform probabilistic inference where it receives observations for a number of the nodes and then determines whether the user would like to visit the business web site 1506 or the travel web site 1508

-76-

based on the received observations. One skilled in the art will appreciate that the Bayesian network 1500 is merely exemplary and that the Bayesian network used by the exemplary embodiment may have many more nodes.

5

Figure 25B depicts a decision graph 1520 suitable for use in the business node 1506 of the Bayesian network 1500 of Figure 25A. In the decision graph 1520, age vertex 1522 serves as the root vertex of the data structure, sex vertices 1524 and 1526 serve as the intermediate vertices of the data structure, and vertices 1528-1532 serve as the leaf vertices of the data structure, which contain the probabilities for the business vertex 1506 of the Bayesian network 1500. It should be noted that vertex 1530 reflects an equivalence relationship where the probability of a female of age bracket 2 likely visiting business-related web sites and the probability of males of age brackets 0, 1, or 3 likely visiting business-related web sites are equivalent. The process of creating a decision graph for a node in the Bayesian network of an exemplary embodiment is described in detail below.

25 An alternative embodiment of the present invention allows for the possible introduction of cycles into the Bayesian network. The introduction of cycles into the Bayesian network destroys the acyclic property of the Bayesian network so that it is more  
30 appropriately referred to as a cyclic directed graph

-77-

(or a cyclic Bayesian network). Figure 25C depicts a cyclic directed graph 1534, which is similar to Bayesian network 1500, except that cycles have been introduced. Introducing cycles into a Bayesian network is beneficial, because the resulting structure becomes more flexible so that it can more accurately reflect relationships in the data. That is, by enforcing the acyclic nature of a Bayesian network, relationships such as a dual dependency relationship cannot be expressed. For example, with respect to the cyclic directed graph 1534 of Figure 25C, the business node 1506 influences the value of the travel node 1508, and the value of the travel node 1508 influences the value of the business node 1506. Such flexibility provides a more efficient Bayesian network that more accurately reflects the data. Although all arcs 1536-1544 are shown as being bidirectional, one skilled in the art will appreciate that some arcs may be unidirectional.

#### Implementation of Decision Graphs

Referring again now to the computer system of Figure 6, the system is adapted to further include a web analyzer 1614 that utilizes the MBN 508 to perform probabilistic inference and determines whether a given user would like to visit a particular category of web sites. The expert knowledge 506 provided by the expert includes an equivalent sample size. The permanent storage 306 also holds the ECMSS\_c summarizing cases

-78-

reflecting real-world instances of whether a number of users visited business-related or travel-related web sites.

5                   Figure 26A depicts a flowchart of the steps performed by the MBN generator 502 (hereinafter referred to as a Bayesian network generator). At the completion of the Bayesian network generator's processing, a hypothesis-specific Bayesian network  
10                   similar to Bayesian network 1500 of Figure 25A is generated and the nodes of the Bayesian network have a decision graph similar to the decision graph 1520.

                  The first step (step 1802) performed by the  
15                   Bayesian network generator is to initialize the decision graphs corresponding to each node in the belief network. This is done by creating decision graphs that are equivalent to full tables for the initial hypothesis-specific Bayesian network. The  
20                   Bayesian network generator selects a node in the initial hypothesis-specific Bayesian network (step 1804). After selecting a node in the Bayesian network, the Bayesian network generator inserts the counts and the equivalent sample size into the leaves  
25                   of the decision graph of the node (step 1806). The count for a leaf is the number of times each value of the leaf is observed in the ECMSS\_c (stored in the permanent storage 306) for each value of its parent vertices. To better explain the counts stored in the  
30                   leaf, consider decision graph 1904 of Figure 27A, which



-79-

is an example decision graph for the business node 1506. Leaf 1908 of decision graph 1904 contains two counts: one count indicates the number of times in the database where sex = male and business = yes, and the other count indicates the number of times that sex = male and business = no. Leaf 1912 also contains two counts: one count for the number of times in the database where sex = female, age = 2, and business = yes, and the other count is for the number of times sex = female, age = 2, and business = no. Similarly, leaf 1914 contains two counts: one count for the number of times sex = female, age = 0, 1, or 3, and business = yes, and the other count is for sex = female, age = 0, 1, or 3, and business = no. It should be appreciated that if a leaf could be arrived at through more than one path, such as occurs when an equivalency relationship is reflected by the decision graph, the leaf will have additional counts. Next, the Bayesian network generator makes various adjustments to the decision graph and generates a number of candidate decision graphs (step 1808). This step is further discussed below with respect to Figure 27B.

After generating the candidate decision graphs, the Bayesian network generator selects the candidate decision graph with the best score (step 1810). In this step, the Bayesian network generator generates a score for each decision graph generated in step 1808. This score indicates the goodness of the graph at reflecting the data contained

-80-

in the database. This step is performed by performing the following calculation:

$$\log p(B_i^e \parallel C) + \log \left[ \prod_{a=1}^n \prod_{b \in G_a} \frac{\Gamma\left(\frac{Kt_b}{q_a}\right)}{\Gamma\left(N_{ab} + \frac{Kt_b}{q_a}\right)} \prod_{c=1}^{r_a} \frac{\Gamma\left(N_{abc} + \frac{Kt_b}{r_a q_a}\right)}{\Gamma\left(\frac{Kt_b}{r_a q_a}\right)} \right]$$

5

where: "n" is the total number of nodes in the Bayesian network,  $G_a$  is the set of leaves for the decision graph in node A of the Bayesian network,  $r_a$  is the number of states of node A, and  $q_a$  is the number of configurations of the parents of node A, and  $t_b$  is the number of configurations of the parents of node a corresponding to b. The term " $N_{abc}$ " is the expected number of cases where node "a" has a value "c" and the parents of leaf "b" in the decision graph of node "a" are in a state that leads to leaf "b." The term " $N_{ab}$ " is the sum over "c" of " $N_{abc}$ ." When performing this step, most of the leaves of the decision graph will have the counts already stored from the processing performed in step 1806. However, for those newly generated leaves, created during the processing of step 1808 (discussed below), the counts have not been stored. For these leaves, the Bayesian network generator obtains the counts as described above. After scoring each candidate graph, the Bayesian network generator selects the candidate graph with the best score and stores this graph into the node.

25

-81-

Most candidate graphs (other than the first one generated) reflect a single change to a preexisting candidate decision graph where one or more vertices are added. Therefore, when a preexisting decision graph has already been scored, the exemplary embodiment can optimize the scoring step. The exemplary embodiment optimizes the scoring step by obtaining a partial score by only scoring the added vertices, by adding this partial score by the score of the preexisting decision graph, and by subtracting out the portion of the score for parts of the preexisting decision graph that no longer exist (i.e., any vertices or edges that were removed). Those practiced in the art will recognize that a factorable structure prior is required to perform this step.

Next, the Bayesian network generator determines if there are more nodes in the Bayesian network for processing (step 1812). If there are more nodes in the Bayesian network for processing, processing continues to step 1804. However, if there are no more nodes in the Bayesian network for processing, the Bayesian network generator identifies which node has the graph with the best score (step 1814). In this step, the Bayesian network generator compares the score of the graph selected in step 1810 for each node to determine which of the nodes ("the selected node") has the graph whose relative score has improved the most ("the selected decision graph"). The Bayesian network generator then makes the

-82-

change reflected by the selected decision graph by retaining the selected decision graph (step 1816). In this step, the selected decision graph replaces the current decision graph in the selected node.

5

After replacing the graph, the Bayesian network generator updates the Bayesian network (step 1818). In this step, the Bayesian network generator determines if the change made per the selected decision graph reflects that a relationship between the nodes of the Bayesian network exists, which is not currently reflected by the Bayesian network. To do this, the Bayesian network generator determines if the change reflected by the selected decision graph was either a complete split or a binary split on a node that is not currently a parent of the selected node as reflected in the Bayesian network. Both a complete split and a binary split are discussed below. This test is performed to determine whether the Bayesian network structure needs to be updated. In this situation, a node was added into the selected decision graph for the selected node in the Bayesian network, which indicates that the added node influences the probabilities for the selected node. Since the probabilities of the selected node are influenced by a node that is not currently a parent to the selected node in the Bayesian network, an arc is added from the node to the selected node in the Bayesian network to indicate such a relationship. This addition of an arc may introduce a cycle in the alternative embodiment,

30

-83-

but in the exemplary embodiment, since there are restrictions placed on the conditions under which a split occurs, no cycle is introduced.

5           After updating the Bayesian network, the Bayesian network generator adds the scores for all nodes (i.e., the decision graphs in the nodes) together (step 1820). The Bayesian network generator then compares this score for the Bayesian network against  
10   the most recent Bayesian network generated by the Bayesian network generator to determine if this is the best score yet (step 1822). The Bayesian network generator retains the last Bayesian network that is produced. If the score for the most recent Bayesian  
15   network is the best score yet, processing continues to step 1804 to generate another Bayesian network. However, if the score is not the best yet, then the Bayesian network generator outputs the last generated Bayesian network, which is the Bayesian network with  
20   the highest score (step 1824).

Figure 26B depicts a flowchart of the steps performed by the Bayesian network generator in step 1808 of Figure 26A to generate candidate decision  
25   graphs. The processing of this flowchart is performed on the decision graph of a node ("the identified node") of the Bayesian network identified per step 1804 of Figure 26A. The first step performed by the Bayesian network generator is to select a leaf in the decision  
30   graph of the identified node (step 1840). After

-84-

selecting a leaf, the Bayesian network generator performs a complete split to generate a number of new decision graphs (step 1842). In this step, the Bayesian network generator performs a complete split on all nodes of the Bayesian network that are not descendants of the identified node ("non-descendent nodes"). For example, with respect to the Bayesian network 1500 of Figure 25A, if the identified node is the business node 1506, the non-descendant nodes include the age node 1502 and the sex node 1504, but not the travel node 1508, because the travel node is a descendent of the business node. This limitation is enforced so as to prevent the introduction of cycles into the Bayesian network. However, if an alternative embodiment of the present invention is used where cycles are allowed to be introduced into the Bayesian network, then complete splits are performed on all nodes in the Bayesian network other than the parent of the leaf node. When performing a complete split, the Bayesian network generator selects one of the non-descendent nodes described above and replaces the leaf node in the decision graph with a vertex that corresponds to the selected non-descendent node. Then, new leaves are created which depend from the newly created vertex; one leaf vertex is created for each value of the newly added vertex. For example, if the leaf vertex 1908 of the decision graph 1904 of Figure 27A had a complete split performed on the age node, the resulting decision graph appears in Figure 27B where the leaf 1908 of Figure 27A is

-85-

replaced with age vertex 1918 of Figure 27B and leaves 1920-1926 are created, one for each value of the age vertex (i.e., each state of the age node of the Bayesian network). Each complete split on a particular non-descendent node generates a new decision graph which is stored. To conserve space, an exemplary embodiment stores an identification of the change and not the entire decision graph.

After performing a complete split, the Bayesian network generator performs a binary split if the number of states is greater than two (step 1844). In this step, a binary split is performed on the leaf for all nodes that are not descendants of the identified node as reflected in the Bayesian network and for all values for these non-descendent nodes. As stated above, this restriction is enforced to prevent the addition of cycles into the Bayesian network. However, an alternative embodiment does not enforce this restriction. In a binary split operation, a leaf is replaced with a vertex that corresponds to one of the non-descendant nodes, and two leaves are generated from the newly created vertex node: one of the leaves contains a single value and the other leaf contains all other values. For example, in the decision graph 1904 of Figure 27A, if leaf 1908 had a binary split performed on the age variable, the leaf 108 of Figure 27A would be replaced with age vertex 1930 as shown in Figure 27C and two leaves 1932 and 1934 would be generated for that vertex. The first leaf 932 would

-86-

contain one value (e.g., 1) and the second leaf 1934 would be for all other values of the age vertex 1930 (e.g., 0, 2 and 3). As stated above, the binary splits on the leaf will be performed for all non-descendent nodes and for each value of each non-descendent node. Thus, when a node has n values, a binary split is performed on this node n times. For example, since the age node has four values, four splits would occur: (1) one leaf would have a value of 0, and the other leaf would have a value of 1, 2, or 3; (2) one leaf would have a value of 1, and the other leaf would have a value of 0, 2, or 3; (3) one leaf would have a value of 2, and the other leaf would have a value of 0, 1, or 3; (4) one leaf would have a value of 3, and the other leaf would have a value of 0, 1, or 2. The Bayesian network generator stores identifications of the changes reflected by these binary splits.

After performing a binary split, the Bayesian network generator merges all pairs of leaf nodes together (step 1846). In this step, the Bayesian network generator generates a number of new decision graphs by merging the leaf node selected in step 1840 with each other leaf node to form a single vertex. For example, with respect to the decision graph 1904 of Figure 27A, leaf 1908 and leaf 1912 can be merged into a single leaf 1938 as depicted in Figure 27D. After merging all pairs of leaf nodes, the Bayesian network generator determines if the decision graph has more leaves for processing. If so, processing continues to



-87-

step 1840. Otherwise, processing ends. Although the exemplary embodiment is described as performing a complete split, a binary split, and a merge, one skilled in the art will appreciate that other  
5 operations can be performed.

Figure 28 depicts a flowchart of the steps performed by the web analyzer 1614 (Figure 6) of an exemplary embodiment of the present invention. The web  
10 analyzer first receives the MBN output by the Bayesian network generator (step 2002). After receiving the MBN network, the web analyzer receives a request from a user containing values (step 2004). In this step, the Bayesian network generator receives observations or  
15 values for a number of the nodes of the MBN. For example, the user may input their age and sex. The web analyzer then performs probabilistic inference and ranks the web site categories, business and travel, by the likelihood that the user would like to visit them  
20 (step 2006). In this step, any standard Bayesian network inference algorithm, such as the one described in Jensen, Lauritzen, and Olesen, "Bayesian Updating in Recursive Graphical Models by Local Computations",  
Technical Report R-89-15, Institute of Electronic  
25 Systems, Aalborg University, Denmark, may be used by an exemplary embodiment of the present invention. Before using such an inference algorithm, the probabilities of each Bayesian network node is expressed as a table. Such an inference algorithm and its usage is described  
30 in greater detail in U.S. Patent Application

-88-

No. 08/602,238, entitled "Collaborative Filtering Utilizing a Belief Network," which has previously been incorporated by reference. If the Bayesian network of an alternative embodiment is used, where the Bayesian network contains cycles, the inference algorithm used is to merely access the decision graph with the values for the nodes received in step 2004 to determine the probability. In this situation, all parent nodes of a node for which inference is requested should have a value provided. After performing probabilistic inference and ranking the nodes reflecting categories of web sites, the web analyzer determines if there are more requests from the user (step 2008). If there are more requests, processing continues to step 2004. However, if there are no more requests, processing ends.

Using a Mixture of Bayesian Networks to Perform Collaborative Filtering:

Collaborative filtering has been described in the above-reference application entitled "Collaborative Filtering Utilizing A Belief Network". The mixture of Bayesian networks of the present invention can be employed to carry out the same type of collaborative filtering in a more powerful way. In this case, the collaborative filtering described in the above-referenced application is a special limited case of a collaborative filter using the present invention, a mixture of Bayesian networks. In the special limited case of the prior above-referenced application,

-89-

there are no arcs in the HSBNs, there are no hidden variables in the HSBNs and there is no structure search step (block 38 of Figure 12). Thus, the present invention provides a more general and more powerful network for collaborative filtering. Such collaborative filtering is carried out by an appropriate assignment of variables of the mixture of Bayesian networks already described herein. The following is a detailed description of how to assign those variables in order to carry out collaborative filtering using the embodiments of the present invention.

Figure 29 depicts an exemplary typical HSBN 2400 within an MBN utilized to determine preferences of a user for a television show. In the exemplary embodiment, Bayesian networks are implemented as an acyclic directed graph data structure with the variables in the Bayesian network corresponding to nodes in the data structure. The Bayesian network 2400 contains a number of variables (or nodes) 2402, 2404, 2406, 2408, 2410, 2412, 2414 and 2416. Two of these variables, 2402 and 2404, reflect causal attributes and are sometimes referred to as causal variables. A "causal attribute" is an attribute that has a causal effect on caused attributes. The caused attributes in the Bayesian network 2400 are reflected by variables 2406, 2408, 2410, 2412 and 2414. These variables are known as caused attributes (or caused variables) because their value is causally influenced

-90-

by the causal variables. Caused attributes can be of two types: preference or non-preference. Preference caused attributes contain the preferences to be predicted. Non-preference caused attributes are

5 causally influenced by the causal attributes, but are not preferences because the system is not used to predict their value. Non-preference caused attributes are further discussed below. For example, variable 2414 is a preference caused attribute

10 indicating whether a particular user likes the "Power Rangers" television show and variable 2402 is a causal attribute whose value has a causal effect on variable 2414. That is, since "Power Rangers" is primarily enjoyed by children, the younger the age

15 variable, the more likely it is that the user will enjoy the "Power Rangers" show.

As part of the prior knowledge, an administrator also supplies a prior probability that

20 indicates the administrator's level of confidence that the Bayesian network adequately predicts the preferences of the user and a range of a number of states for any hidden variables in the Bayesian network 2400. For example, the administrator may

25 indicate that the hidden variable 2416 contains between five and ten states based on their own knowledge. Each of these states corresponds to a cluster of users in the database that have similar preferences. The exemplary embodiment during its processing will

30 determine which number of these states most accurately

-91-

reflects the data in the database 316. In other words, the exemplary embodiment will determine a number within the range that is the best grouping of clusters in the database as described above in this specification.

5

While the present invention has been described with reference to a exemplary embodiment thereof, those skilled in the art will know of various changes in form that may be made without departing from the spirit and scope of the claimed invention as defined in the appended claims. Such changes may include parallelization of some of the computations described herein or the use of other probability distributions.

10

-92-

What is claimed is:

- 1        1.    A method of using a set of observed data for  
2        improving the structure and parameters of a mixture of  
3        Bayesian networks comprising plural hypothesis-specific  
4        Bayesian networks (HSBNs) having nodes corresponding to  
5        hidden and observed variables, each of said nodes  
6        storing a set of parameters and structure representing  
7        dependence relationships among said nodes, said method  
8        comprising:  
9                choosing a number of said HSBNs;  
10                choosing a number of states of said discrete  
11        variables, and initializing said HSBNs;  
12                for each one of said HSBNs conducting a  
13        parameter search for a set of changes in said parameters  
14        which improves the goodness of said one HSBN in  
15        predicting said observed data, and modifying the  
16        parameters of said one HSBN accordingly;  
17                for each one of said HSBNs, computing a  
18        structure score of said one HSBN reflecting the goodness  
19        of said one HSBN in predicting said observed data,  
20        conducting a structure search for a change in said  
21        structure which improves said structure search score,  
22        and modifying the structure of said one HSBN  
23        accordingly.
- 1        2.    The method of claim 1 wherein the step of computing  
2        a structure score of said one HSBN comprises:  
3                computing from said observed data expected  
4        complete model sufficient statistics (ECMSS);

-93-

5           computing from said ECMSS sufficient  
6   statistics for said one HSBN;  
7           computing said structure score from said  
8   sufficient statistics.

1   3.   The method of claim 2 wherein the step of computing  
2   said ECMSS comprises:

3           computing the probability of each combination  
4   of the states of the discrete hidden and observed  
5   variables;

6           forming a vector for each observed case in  
7   said set of observed data, each entry in said vector  
8   corresponding to a particular one of the combinations of  
9   the states of said discrete variables; and

10          summing the vectors over plural cases of said  
11   observed data.

1   4.   The method of claim 3 wherein the step of forming a  
2   vector is such that each entry in said vector is formed  
3   to have plural sub-entries comprising:

4           (a) the probability of the one combination of  
5   the states of the discrete variables,

6           (b) sub-entry vectors representing the states  
7   of the continuous variables.

1   5.   The method of claim 4 wherein the probability of  
2   the one combination of the states of the discrete  
3   variables is computed by inference in said mixture of  
4   Bayesian networks.

1       6.    The method of claim 4 wherein each sub-entry is  
2       formed such that said sub-entry vector has a vector  
3       multiplier corresponding to the probability of the one  
4       combination of the states of the discrete variables.

1       7.    The method of claim 6 wherein the step of computing  
2       sufficient statistics from said ECMSS comprises  
3       computing from said ECMSS the following:  
4               (a) mean,  
5               (b) scatter,  
6               (c) sample size.

1       8.    The method of claim 1 wherein the steps of  
2       conducting a parameter search and modifying said  
3       parameters are repeated consecutively until a parameter  
4       search convergence criteria is met.

1       9.    The method of claim 1 further comprising:  
2               repeating the steps of conducting a parameter  
3       search, computing the structure score and conducting a  
4       structure search until a structure search convergence  
5       criteria is met.

1       10.   The method of claim 8 further comprising:  
2               repeating the steps of conducting a parameter  
3       search, computing the structure score and conducting a  
4       structure search until a structure search convergence  
5       criteria is met.



-95-

1 11. The method of claim 9 wherein said parameter search  
2 convergence criteria is a determination of whether the  
3 parameter search has converged at a local optimum.

1 12. The method of claim 9 wherein said parameter search  
2 convergence criteria is a determination of whether the  
3 parameter search has been repeated a certain number of  
4 times.

1 13. The method of claim 12 wherein said certain number  
2 of times is a set number.

1 14. The method of claim 12 wherein said certain number  
2 of times is a function of the number of times the  
3 structure search has been repeated.

1 15. The method of claim 12 wherein said parameter  
2 search convergence criteria limits the repetition of  
3 said parameter search to a limited number of repetitions  
4 and wherein said parameter search is repeated after  
5 convergence of said structure search.

1 16. The method of claim 10 wherein said structure  
2 search convergence criteria comprises a determination of  
3 whether the structure score has worsened since a prior  
4 repetition of said structure search step.

1 17. The method of claim 10 wherein said structure  
2 search criteria comprises a determination of whether a  
3 current performance of the structure search has changed  
4 any of said structure in the one HSBN.

1 18. The method of claim 1 wherein the step of  
2 conducting a structure search comprises:  
3 attempting different modifications to said  
4 structure at each node of said one HSBN;  
5 for each one of said different modifications,  
6 computing the structure score of the one HSBN;  
7 saving those modifications providing  
8 improvements to said structure score.

1 19. The method of claim 1 further comprising computing  
2 a combined score of said mixture of Bayesian networks  
3 from the structure scores of the individual HSBNs.

1 20. The method of claim 19 further comprising  
2 associating said mixture of Bayesian networks with said  
3 combined score.

1 21. The method of claim 20 further comprising choosing  
2 a different number of states of said discrete hidden and  
3 observed variables and repeating said parameter and  
4 structure search steps, to generate a different mixtures  
5 of Bayesian networks and scores thereof for different  
6 numbers of states of said discrete variables.

-97-

1 22. The method of claim 21 further comprising choosing  
2 one the mixture of Bayesian networks having the highest  
3 score.

1 23. The method of claim 21 further comprising weighting  
2 inference outputs of the different mixtures of Bayesian  
3 networks in accordance with their individual scores.

1 24. The method of claim 1 wherein said parameter search  
2 is repeated whenever a performance of said structure  
3 search results in a change in the structure of said  
4 HSNB.

1 25. The method of claim 24 wherein the parameter search  
2 is repeated a limited number of times while the  
3 structure search is always carried out to convergence.

1 26. The method of claim 24 wherein the parameter search  
2 is repeated to convergence and thereafter the structure  
3 search is repeated to convergence.

1 27. The method of claim 24 wherein the parameter search  
2 is repeated by a number of times which is a function of  
3 the number of times the structure search as been  
4 repeated.

1 28. The method of claim 24 wherein said parameter  
2 search is repeated a fixed number of times and said  
3 structure search is repeated a fixed number of times.

1 29. The method of claim 24 wherein the parameter search  
2 is repeated to convergence while the structure search is  
3 repeated a limited number of times.

1 30. The method of claim 24 wherein said parameter  
2 search is repeated a number of times which is a function  
3 of the number of structure searches performed thus far,  
4 while the structure search is repeated a fixed number of  
5 times.

1 31. The method of claim 1 further comprising repeating  
2 the steps of performing said parameter search and said  
3 structure search and interleaving repetitions of said  
4 parameter search and said structure search.

1 32. The method of claim 1 wherein the step of  
2 initializing said HSBNs comprises, for each HSBN:  
3 defining a structural link from each discrete  
4 hidden variable node to each observed variable node and  
5 from each continuous hidden variable node to each  
6 continuous observed variable node;  
7 initialize the parameters in each node.

1 33. The method of claim 32 wherein the step of  
2 initializing the parameters employs the same initial  
3 parameters from node to node.

-99-

1       34. The method of claim 32 wherein the step of  
2       initializing the parameters comprises:  
3               removing hidden nodes and adjacent arcs from  
4       the HSBN;  
5               determining the maximum *a posteriori* (MAP)  
6       configuration of said HSBN given the training data;  
7               creating a conjugate distribution for said MAP  
8       parameters;  
9               for each observed node in said HSBN and for  
10       each MAP configuration of the observed node's parents,  
11       initializing the parameters of the local distribution  
12       family of said observed node from said conjugate  
13       distribution;  
14               for each hidden discrete node in said HSBN and  
15       for each MAP configuration of said hidden discrete  
16       node's parents, if any, initialize the parameters of the  
17       local distribution family of said hidden discrete node  
18       to be a fixed distribution.

1       35. The method of claim 34 wherein said HSBN contains  
2       no hidden continuous variables.

1       36. The method of claim 32 wherein the step of  
2       initializing the parameters comprises:  
3               initializing said parameters randomly.

1       37. The method of claim 36 wherein said HSBN contains  
2       at least one hidden continuous variable.

-100-

1       38. The method of claim 36 wherein the step of  
2       initializing said parameters randomly comprises one of:  
3               (a) setting the parameters of said HSBN to be  
4       equal;  
5               (b) drawing the parameters from a Dirichlet  
6       distribution.

1       39. The method of claim 1 wherein the step of  
2       performing the parameter search comprises searching for  
3       a change in the parameters in each node which improves  
4       the performance of said one HSBN in predicting said  
5       observed data.

1       40. The method of claim 1 wherein one of said hidden  
2       variables is a common external discrete hidden variable  
3       not represented by any node in said mixture of Bayesian  
4       networks, and wherein the number of HSBNs in said  
5       mixture of Bayesian networks is equal to the number of  
6       states of said common external discrete hidden variable.

1       41. A method for finding the likeliest number of states  
2       of hidden discrete variables in a mixture of Bayesian  
3       networks comprising plural hypothesis-specific Bayesian  
4       networks (HSBNs) having nodes corresponding to hidden  
5       and observed variables, each of said nodes storing a  
6       structure and a set of parameters representing causal  
7       relationships among said nodes, said method of training  
8       comprising:

-101-

9           choosing successive numbers of states of said  
10 discrete hidden and observed variables, and for each one  
11 of said successive numbers of states:  
12                 initializing said HSBNs;  
13                 for each one of said HSBNs conducting a  
14 parameter search for a set of changes in said parameters  
15 which improves the goodness of said one HSBN in  
16 predicting said observed data, and modifying the  
17 parameters of said one HSBN accordingly;  
18                 for each one of said HSBNs, computing a  
19 structure score of said one HSBN reflecting the goodness  
20 of said one HSBN in predicting said observed data,  
21 conducting a structure search for a change in said  
22 structure which improves said structure search score,  
23 and modifying the structure of said one HSBN  
24 accordingly;  
25                 computing a combined score of the mixture of  
26 Bayesian networks corresponding to the current number of  
27 states of said discrete variables; and  
28                 choosing the mixture of Bayesian networks having  
29 the best score.

1       42. The method of claim 41 wherein one of said hidden  
2 variables is a common external discrete hidden variable  
3 not represented by any node in said mixture of Bayesian  
4 networks, and wherein the number of HSBNs in said  
5 mixture of Bayesian networks is equal to the number of  
6 states of said common external discrete hidden variable,  
7 whereby each mixture of Bayesian networks corresponds to

-102-

8 an assumption of a certain number of states of said  
9 common external discrete hidden variable.

1 43. The method of claim 41 wherein the step of  
2 computing a structure score of said one HSBN comprises:  
3 computing from said observed data expected  
4 complete model sufficient statistics (ECMSS);  
5 computing from said ECMSS sufficient  
6 statistics for said one HSBN;  
7 computing said structure score from said  
8 sufficient statistics.

1 44. The method of claim 43 wherein the step of  
2 computing said ECMSS comprises:  
3 computing the probability of each combination  
4 of the states of the discrete hidden and observed  
5 variables;  
6 forming a vector for each observed case in  
7 said set of observed data, each entry in said vector  
8 corresponding to a particular one of the combinations of  
9 the states of said discrete variables; and  
10 summing the vectors over plural cases of said  
11 observed data.

1 45. The method of claim 44 wherein the step of forming  
2 a vector is such that each entry in said vector is  
3 formed to have plural sub-entries comprising:  
4 (a) the probability of the one combination of  
5 the states of the discrete variables,



6 (b) sub-entry vectors representing the states  
7 of the continuous variables.

1 46. The method of claim 45 wherein the probability of  
2 the one combination of the states of the discrete  
3 variables is computed by inference in said mixture of  
4 Bayesian networks.

1 47. The method of claim 45 wherein each sub-entry is  
2 formed such that said sub-entry vector has a vector  
3 multiplier corresponding to the probability of the one  
4 combination of the states of the discrete variables.

1 48. The method of claim 43 wherein the step of  
2 computing sufficient statistics from said ECMSS  
3 comprises computing from said ECMSS the following:

- 4 (a) mean,
- 5 (b) scatter,
- 6 (c) sample size.

1 49. The method of claim 41 wherein the steps of  
2 conducting a parameter search and modifying said  
3 parameters are repeated consecutively until a parameter  
4 search convergence criteria is met.

1 50. The method of claim 41 further comprising:  
2 repeating the steps of conducting a parameter  
3 search, computing the structure score and conducting a  
4 structure search until a structure search convergence  
5 criteria is met.

-104-

1 51. The method of claim 49 further comprising:  
2 repeating the steps of conducting a parameter  
3 search, computing the structure score and conducting a  
4 structure search until a structure search convergence  
5 criteria is met.

1 52. The method of claim 50 wherein said parameter  
2 search convergence criteria is a determination of  
3 whether the parameter search has converged at a local  
4 optimum.

1 53. The method of claim 50 wherein said parameter  
2 search convergence criteria is a determination of  
3 whether the parameter search has been repeated a certain  
4 number of times.

1 54. The method of claim 53 wherein said certain number  
2 of times is a set number.

1 55. The method of claim 53 wherein said certain number  
2 of times is a function of the number of times the  
3 structure search has been repeated.

1 56. The method of claim 53 wherein said parameter  
2 search convergence criteria limits the repetition of  
3 said parameter search to a limited number of repetitions  
4 and wherein said parameter search is repeated after  
5 convergence of said structure search.

-105-

1       57. The method of claim 51 wherein said structure  
2       search convergence criteria comprises a determination of  
3       whether the structure score has worsened since a prior  
4       repetition of said structure search step.

1       58. The method of claim 51 wherein said structure  
2       search criteria comprises a determination of whether a  
3       current performance of the structure search has changed  
4       any of said structure in the one HSBN.

1       59. The method of claim 41 wherein the step of  
2       conducting a structure search comprises:  
3               attempting different modifications to said  
4       structure at each node of said one HSBN;  
5               for each one of said different modifications,  
6       computing the structure score of the one HSBN;  
7               saving those modifications providing  
8       improvements to said structure score.

1       60. The method of claim 41 further comprising computing  
2       a combined score of said mixture of Bayesian networks  
3       from the structure scores of the individual HSBNs.

1       61. The method of claim 60 further comprising  
2       associating said mixture of Bayesian networks with said  
3       combined score.

1       62. The method of claim 61 further comprising choosing  
2       a different number of states of said discrete hidden and  
3       observed variables and repeating said parameter and

-106-

4 structure search steps, to generate a different mixtures  
5 of Bayesian networks and scores thereof for different  
6 numbers of states of said discrete variables.

1 63. The method of claim 62 further comprising choosing  
2 one the mixture of Bayesian networks having the highest  
3 score.

1 64. The method of claim 62 further comprising weighting  
2 inference outputs of the different mixtures of Bayesian  
3 networks in accordance with their individual scores.

1 65. The method of claim 41 wherein said parameter  
2 search is repeated whenever a performance of said  
3 structure search results in a change in the structure of  
4 said HSN.

1 66. The method of claim 65 wherein the parameter search  
2 is repeated a limited number of times while the  
3 structure search is always carried out to convergence.

1 67. The method of claim 65 wherein the parameter search  
2 is repeated to convergence and thereafter the structure  
3 search is repeated to convergence.

1 68. The method of claim 65 wherein the parameter search  
2 is repeated by a number of times which is a function of  
3 the number of times the structure search as been  
4 repeated.

-107-

1       69. The method of claim 65 wherein said parameter  
2       search is repeated a fixed number of times and said  
3       structure search is repeated a fixed number of times.

1       70. The method of claim 65 wherein the parameter search  
2       is repeated to convergence while the structure search is  
3       repeated a limited number of times.

1       71. The method of claim 65 wherein said parameter  
2       search is repeated a number of times which is a function  
3       of the number of structure searches performed thus far,  
4       while the structure search is repeated a fixed number of  
5       times.

1       72. The method of claim 41 further comprising repeating  
2       the steps of performing said parameter search and said  
3       structure search and interleaving repetitions of said  
4       parameter search and said structure search.

1       73. The method of claim 41 wherein the step of  
2       performing the parameter search comprises searching for  
3       a change in the parameters in each node which improves  
4       the performance of said one HSNB in predicting said  
5       observed data.

1       74. A computer-readable medium storing  
2       computer-readable instructions for carrying out the  
3       steps of claim 17.

-108-

1 75. A computer-readable medium storing  
2 computer-readable instructions for carrying out the  
3 steps of claim 41.

1 76. A method converting a set of observed incomplete  
2 data into complete statistics for training a mixture of  
3 Bayesian networks comprising plural hypothesis-specific  
4 Bayesian networks (HSBNs) having nodes corresponding to  
5 hidden and observed variables, each of said nodes  
6 storing a structure and a set of parameters representing  
7 causal relationships among said nodes, said method of  
8 training comprising:

9 choosing a number of states of said discrete  
10 hidden and observed variables, and initializing said  
11 HSBNs;

12 for each one of said HSBNs conducting a  
13 parameter search for a set of changes in said parameters  
14 which improves the goodness of said one HSBN in  
15 predicting said observed data, and modifying the  
16 parameters of said one HSBN accordingly;

17 for each one of said HSBNs, computing a  
18 structure score of said one HSBN reflecting the goodness  
19 of said one HSBN in predicting said observed data,  
20 conducting a structure search for a change in said  
21 structure which improves said structure search score,  
22 and modifying the structure of said one HSBN  
23 accordingly;

24 wherein the step of computing a structure  
25 score of said one HSBN comprises:

-109-

26                    computing from said observed data  
27       expected complete model sufficient statistics (ECMSS),  
28                    computing from said ECMSS sufficient  
29       statistics for said one HSBN,  
30                    computing said structure score from said  
31       sufficient statistics;  
32                    wherein the step of computing said ECMSS  
33       comprises:  
34                    computing the probability of each  
35       combination of the states of the discrete hidden and  
36       observed variables;  
37                    forming a vector for each observed case  
38       in said set of observed data, each entry in said vector  
39       corresponding to a particular one of the combinations of  
40       the states of said discrete variables; and  
41                    summing the vectors over plural cases of  
42       said observed data whereby to render a complete set of  
43       information.

1       77. The method of claim 76 wherein the step of forming  
2       a vector is such that each entry in said vector is  
3       formed to have plural sub-entries comprising:  
4                    (a) the probability of the one combination of  
5       the states of the discrete variables,  
6                    (b) sub-entry vectors representing the states  
7       of the continuous variables.

-110-

1 78. The method of claim 77 wherein the probability of  
2 the one combination of the states of the discrete  
3 variables is computed by inference in said mixture of  
4 Bayesian networks.

1 79. The method of claim 78 wherein each sub-entry is  
2 formed such that said sub-entry vector has a vector  
3 multiplier corresponding to the probability of the one  
4 combination of the states of the discrete variables.

1 80. The method of claim 76 wherein the step of  
2 computing sufficient statistics from said ECMSS  
3 comprises computing from said ECMSS the following:

- 4 (a) mean,
- 5 (b) scatter,
- 6 (c) sample size.

1 81. A computer-readable medium storing a data structure  
2 of expected complete model sufficient statistics (ECMSS)  
3 for a mixture of Bayesian networks comprising plural  
4 hypothesis-specific Bayesian networks (HSBNs) having  
5 nodes corresponding to hidden and observed variables,  
6 each of said nodes storing a structure and a set of  
7 parameters representing causal relationships among said  
8 nodes, said ECMSS data structure formed by the steps of:  
9 computing the probability of each  
10 combination of the states of the discrete hidden and  
11 observed variables;  
12 forming a vector for each observed case  
13 in said set of observed data, each entry in said vector



-111-

14 corresponding to a particular one of the combinations of  
15 the states of said discrete variables; and  
16 summing the vectors over plural cases of  
17 said observed data whereby to render a complete set of  
18 information.

1 82. The computer-readable medium of claim 81 wherein  
2 the step of forming a vector is such that each entry in  
3 said vector is formed to have plural sub-entries  
4 comprising:

5 (a) the probability of the one combination of  
6 the states of the discrete variables,

7 (b) sub-entry vectors representing the states  
8 of the continuous variables.

1 83. The computer-readable medium of claim 82 wherein  
2 the probability of the one combination of the states of  
3 the discrete variables is computed by inference in said  
4 mixture of Bayesian networks.

1 84. The computer-readable medium of claim 83 wherein  
2 each sub-entry is formed such that said sub-entry vector  
3 has a vector multiplier corresponding to the probability  
4 of the one combination of the states of the discrete  
5 variables.

-111-

14 corresponding to a particular one of the combinations of  
15 the states of said discrete variables; and  
16 summing the vectors over plural cases of  
17 said observed data whereby to render a complete set of  
18 information.

1 82. The computer-readable medium of claim 81 wherein  
2 the step of forming a vector is such that each entry in  
3 said vector is formed to have plural sub-entries  
4 comprising:

5 (a) the probability of the one combination of  
6 the states of the discrete variables,

7 (b) sub-entry vectors representing the states  
8 of the continuous variables.

1 83. The computer-readable medium of claim 82 wherein  
2 the probability of the one combination of the states of  
3 the discrete variables is computed by inference in said  
4 mixture of Bayesian networks.

1 84. The computer-readable medium of claim 83 wherein  
2 each sub-entry is formed such that said sub-entry vector  
3 has a vector multiplier corresponding to the probability  
4 of the one combination of the states of the discrete  
5 variables.




FIG. 1




FIG. 3A

320

$p(Z=0   X=0, Y=0)$	.35
$p(Z=0   X=0, Y=1)$	.6
$p(Z=0   X=1, Y=0)$	.6
$p(Z=0   X=1, Y=1)$	.53

FIG. 3B




FIG. 3C




FIG. 2


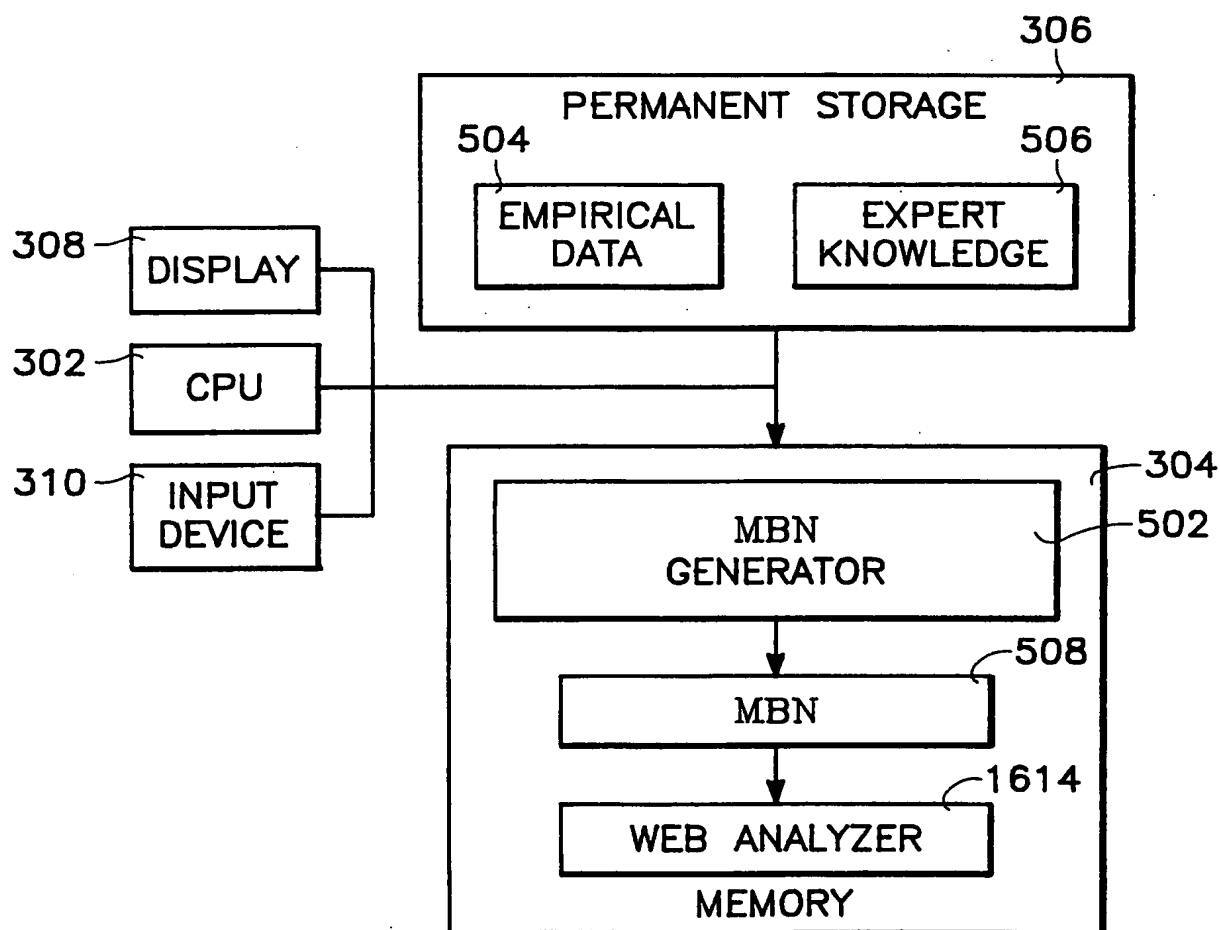
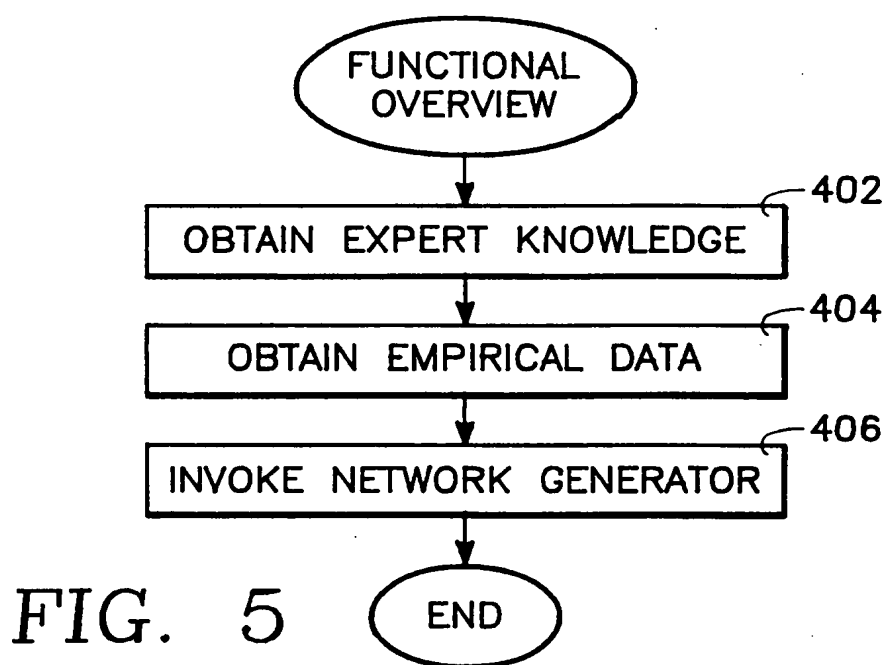


FIG. 4






FIG. 7




FIG. 8




FIG. 9




FIG. 10

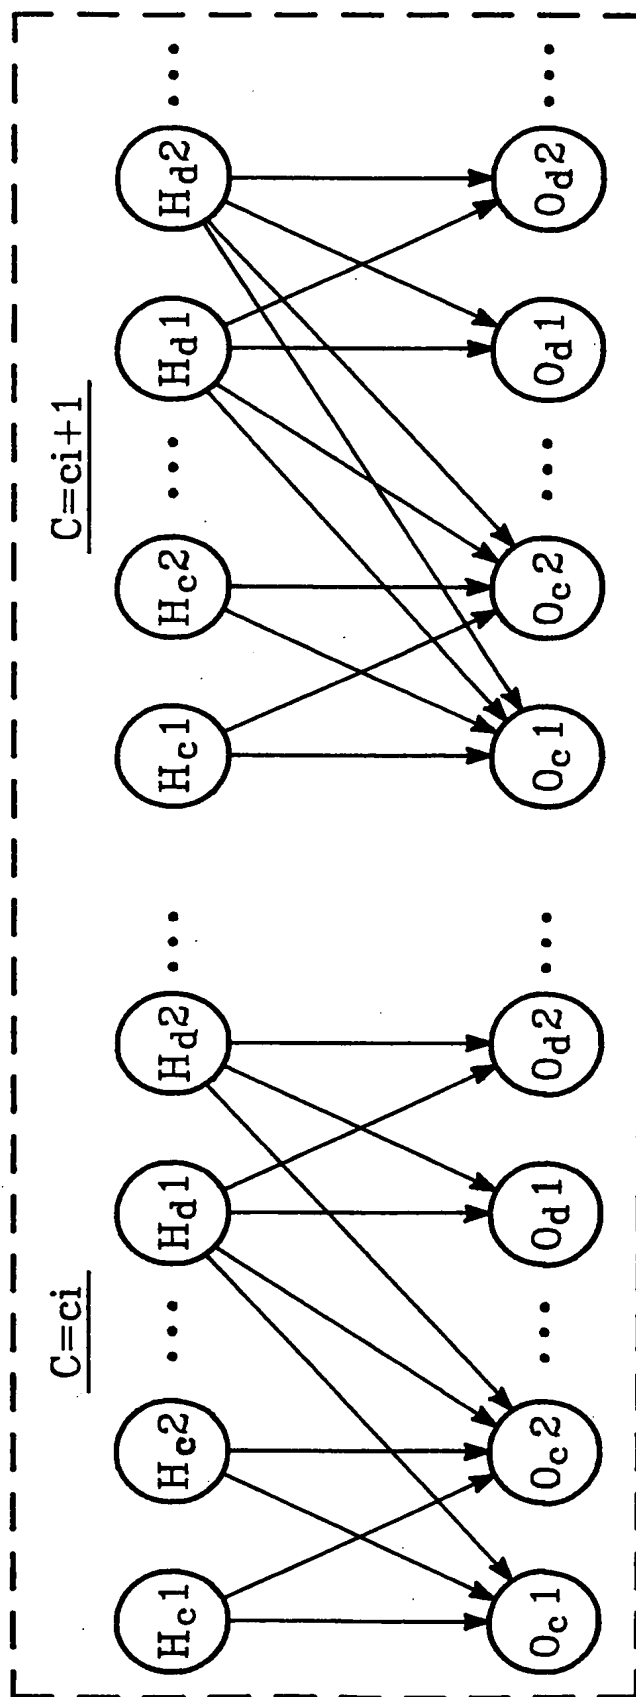


FIG. 11



FIG. 12

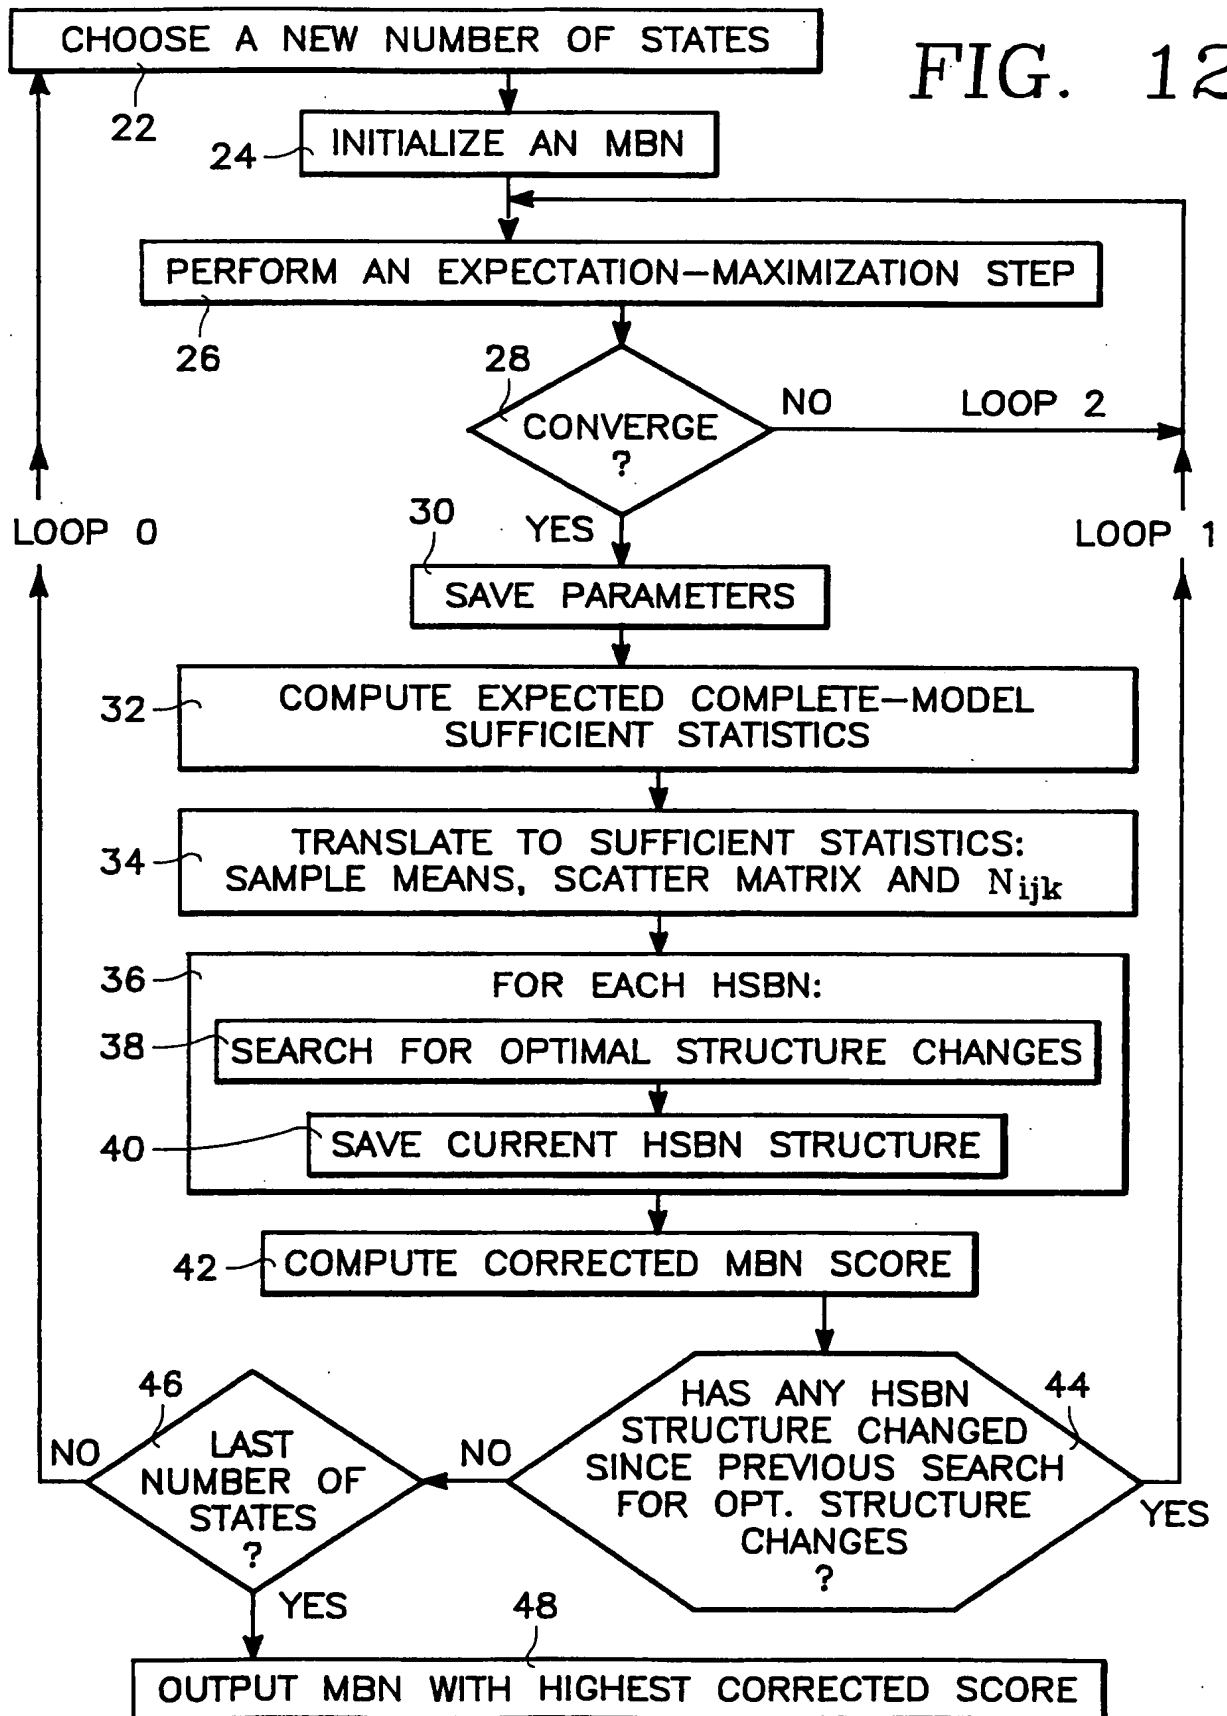


FIG. 13

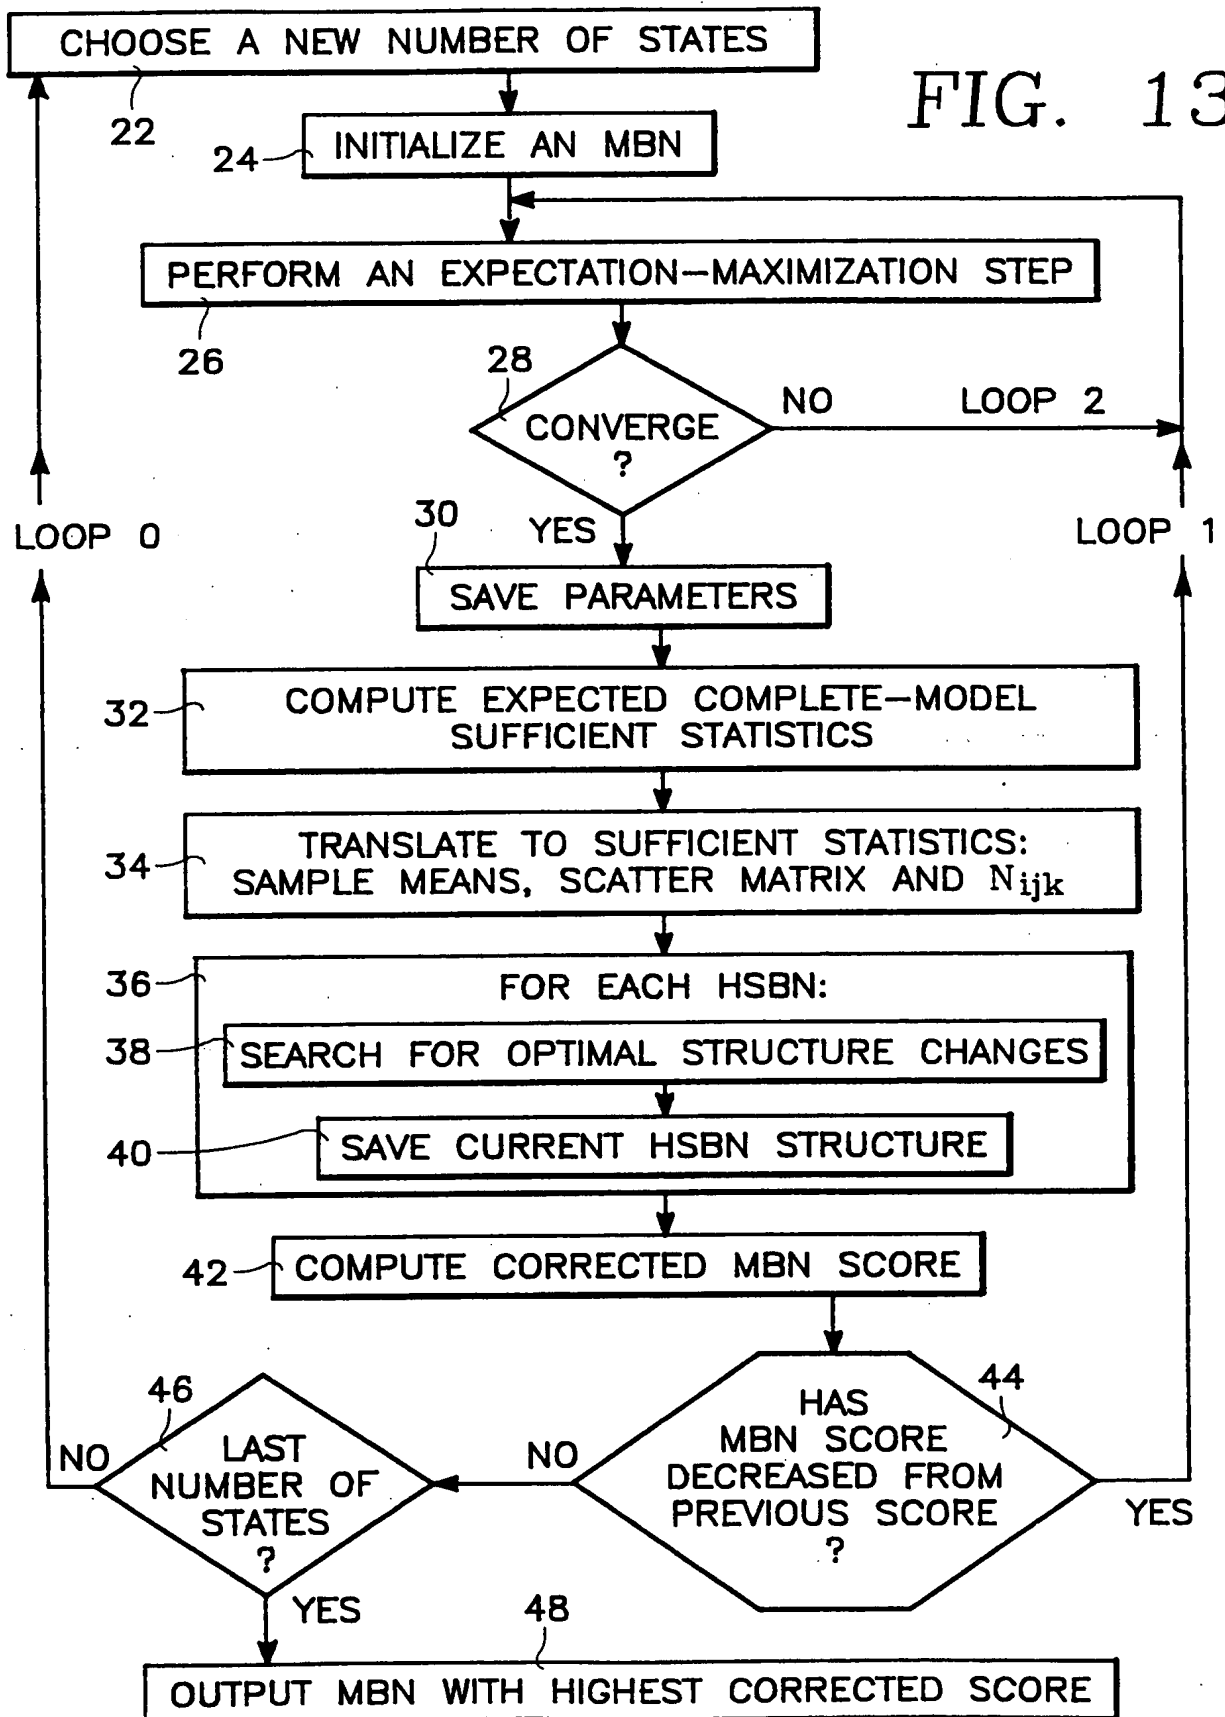


FIG. 14

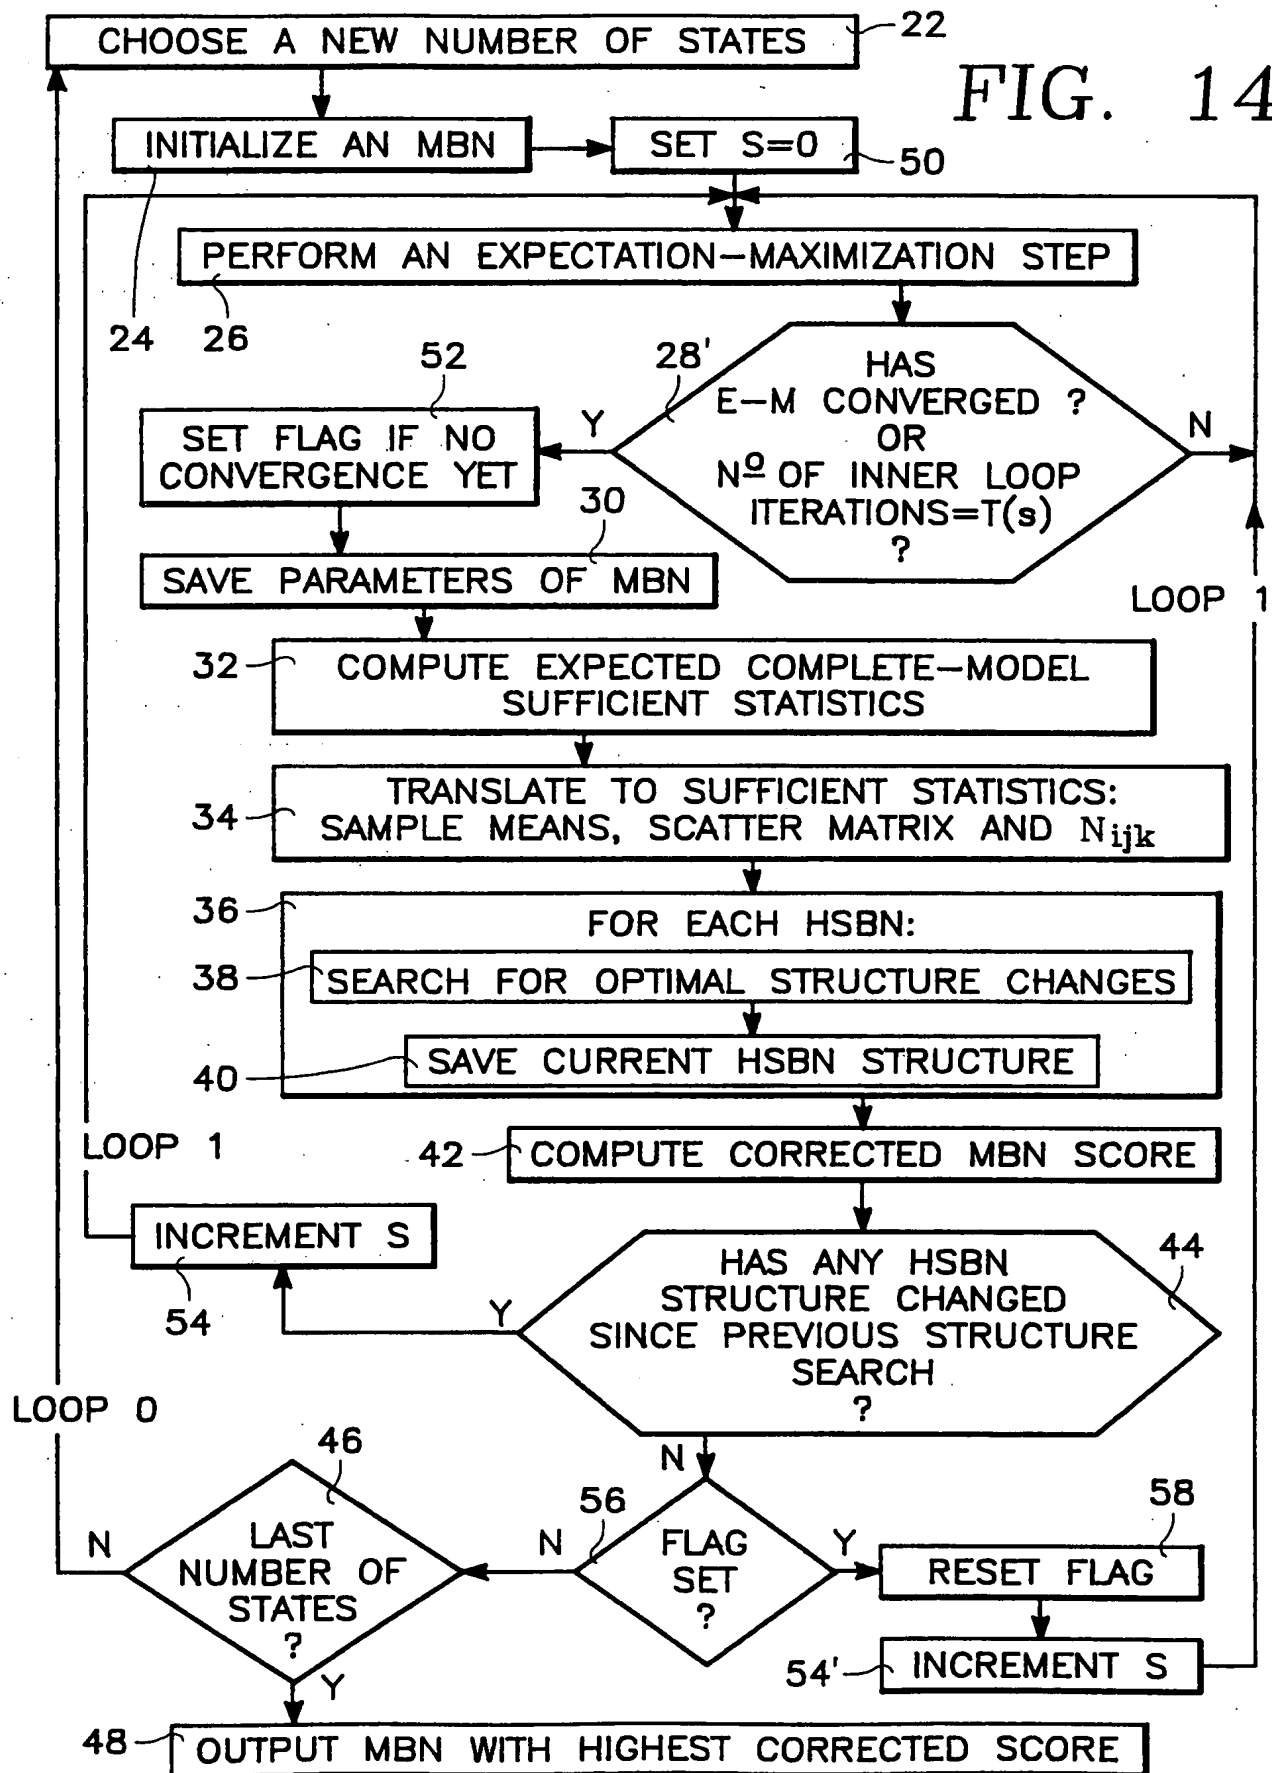
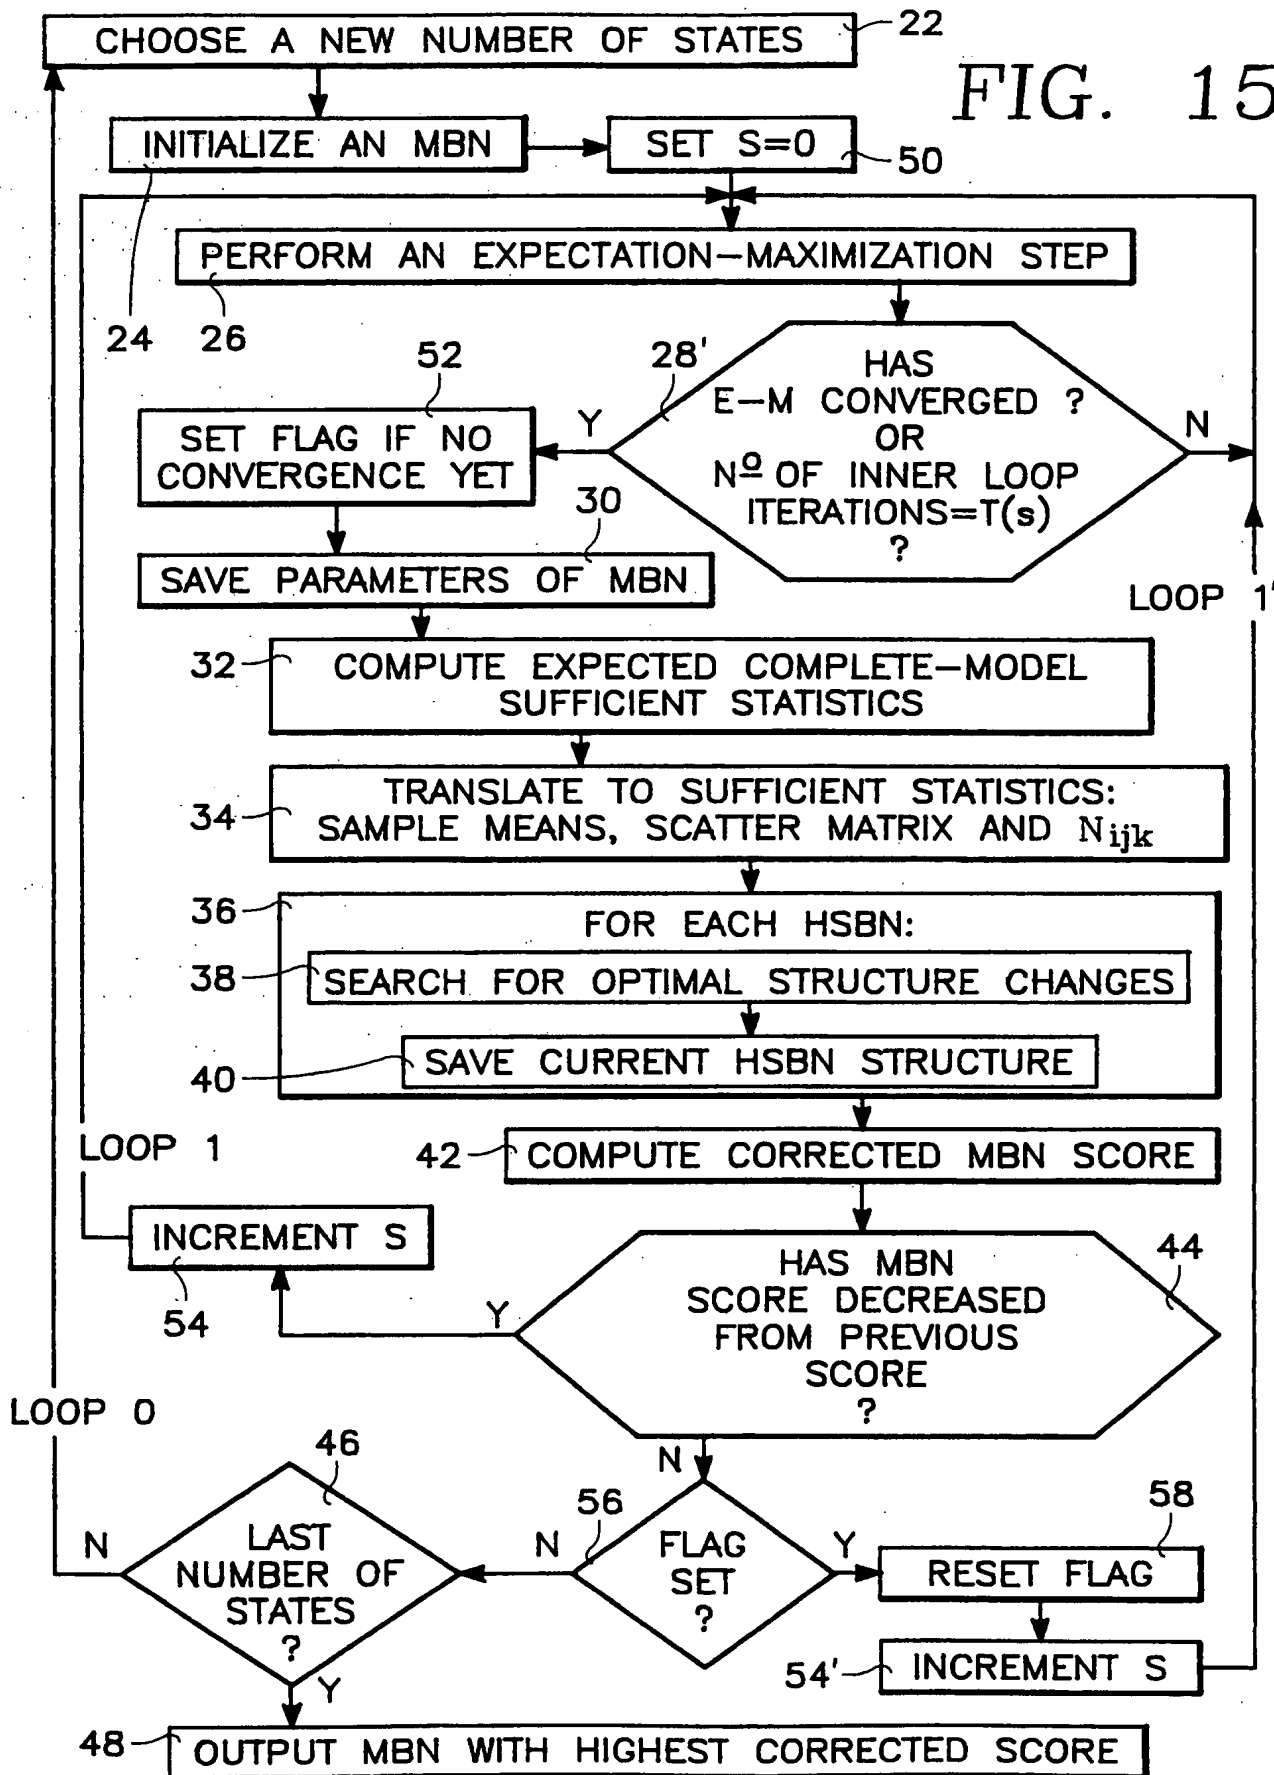


FIG. 15



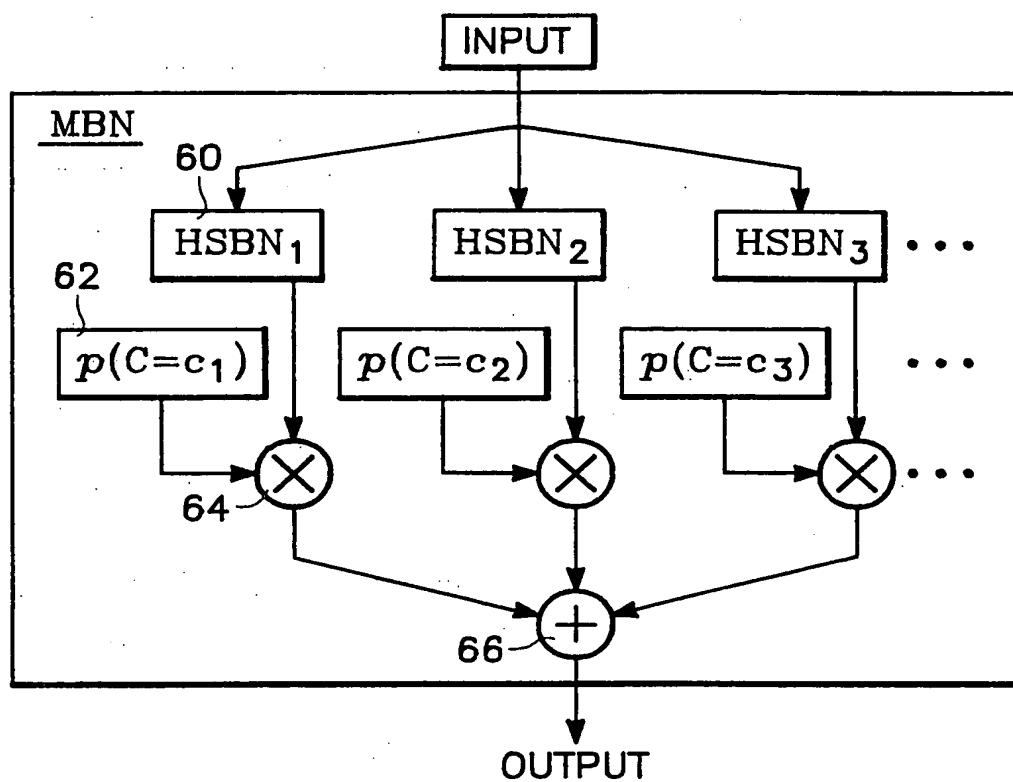


FIG. 16




FIG. 17




FIG. 18

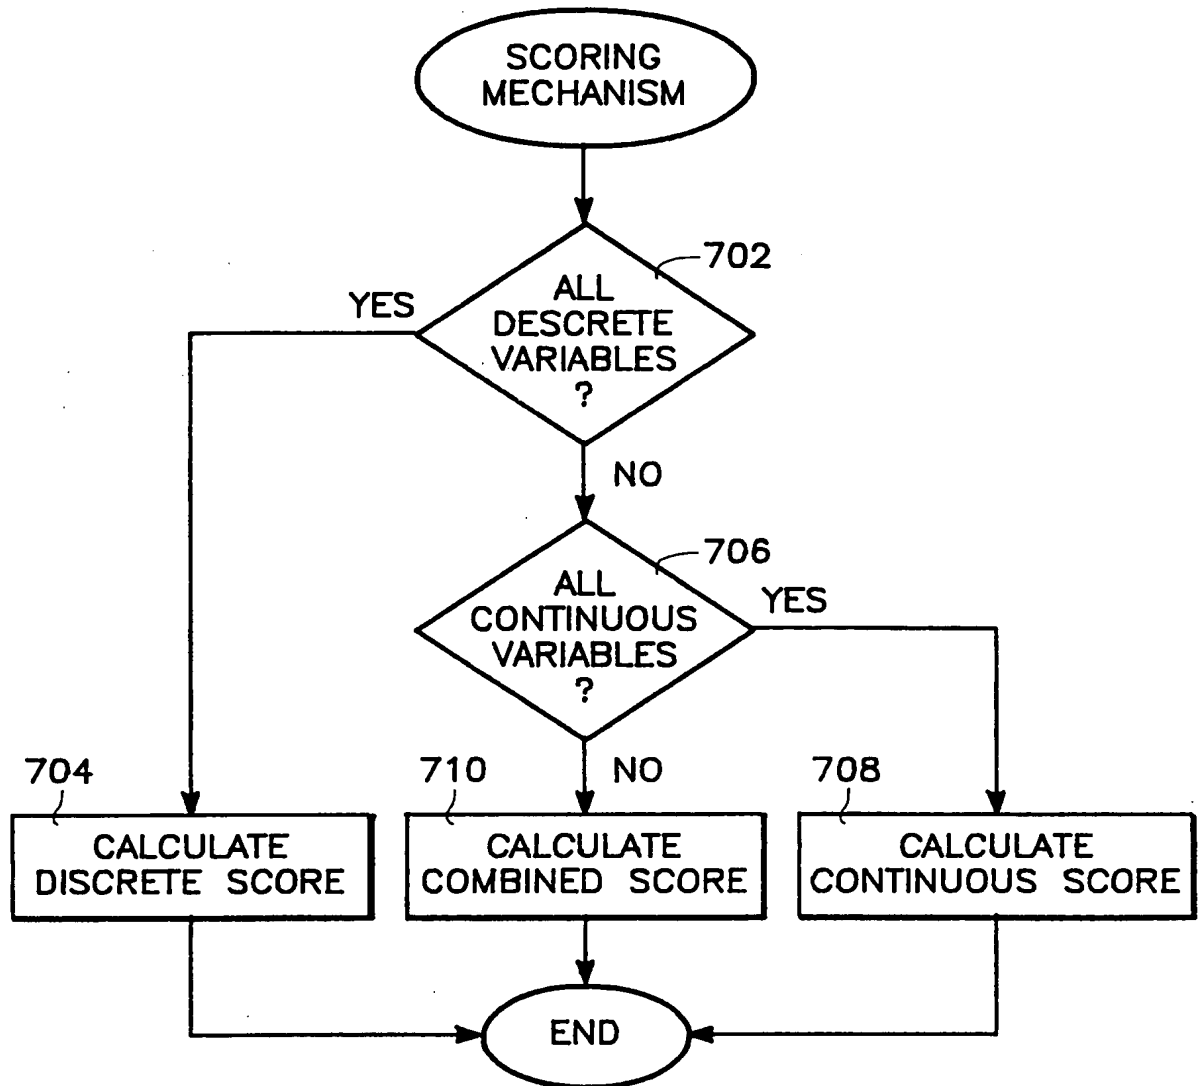


FIG. 19



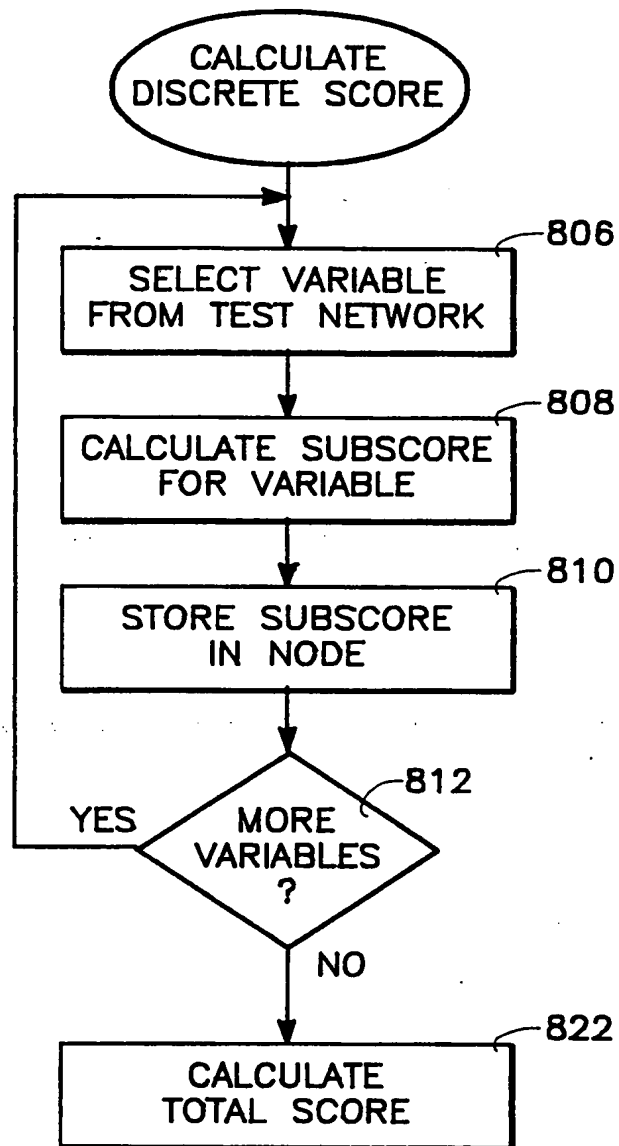


FIG. 20

FIG. 21A




FIG. 21B

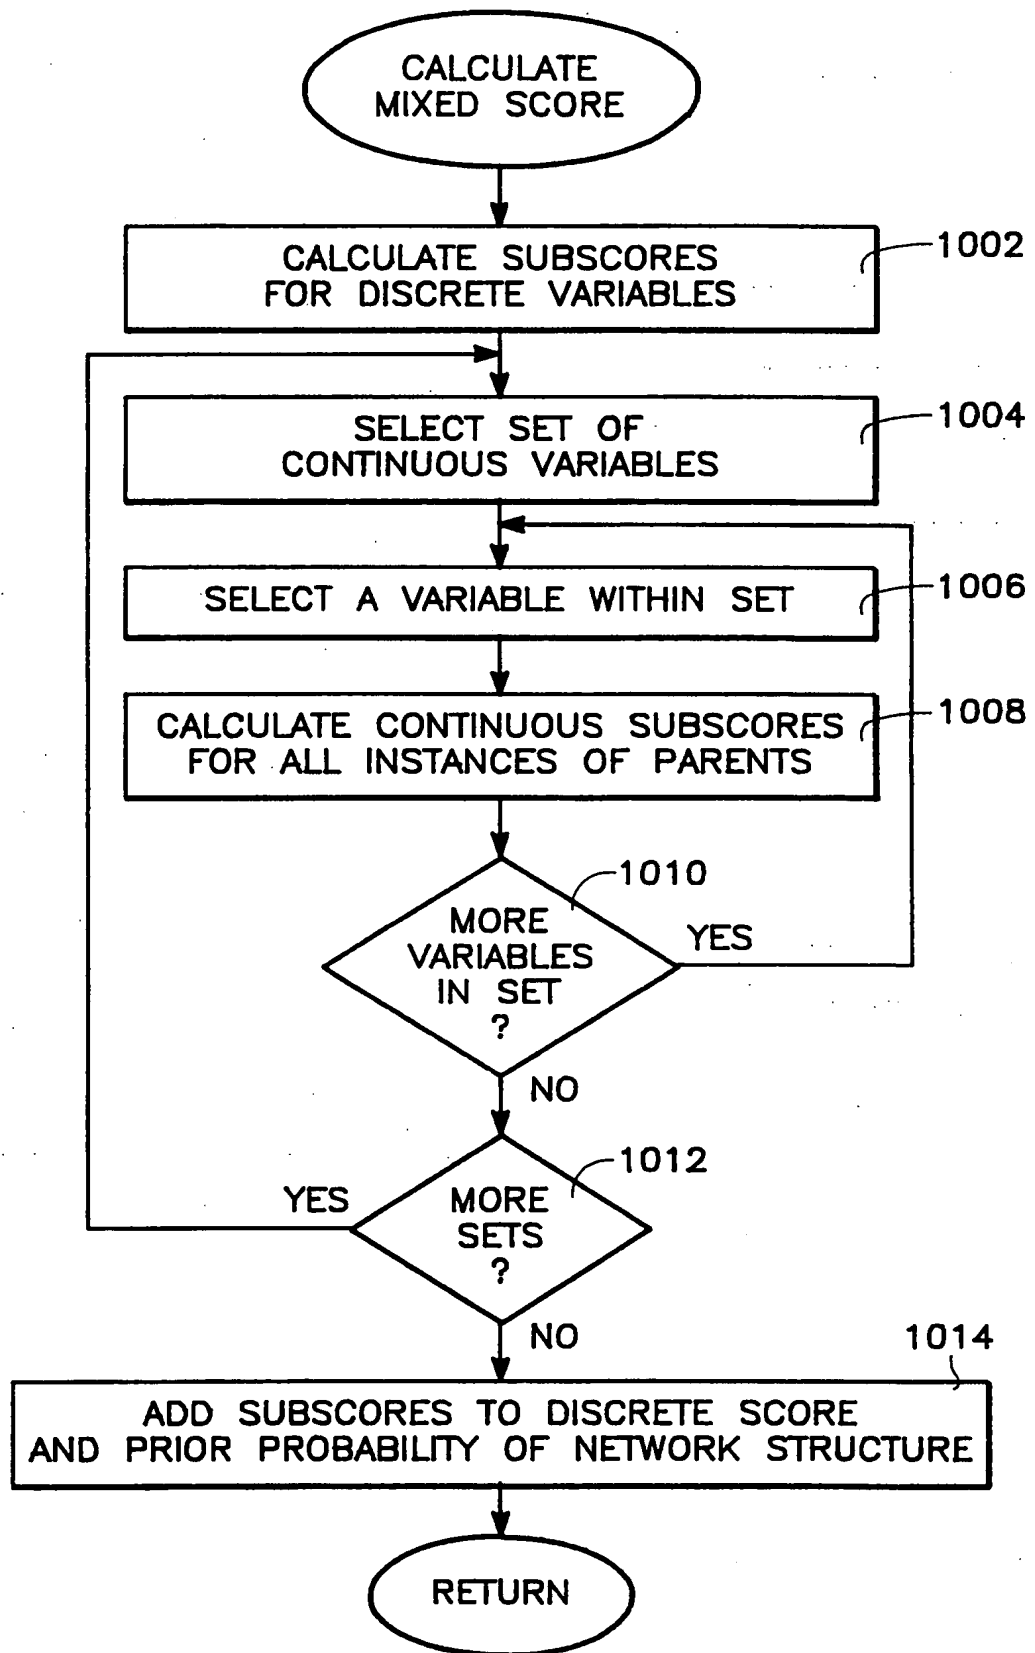

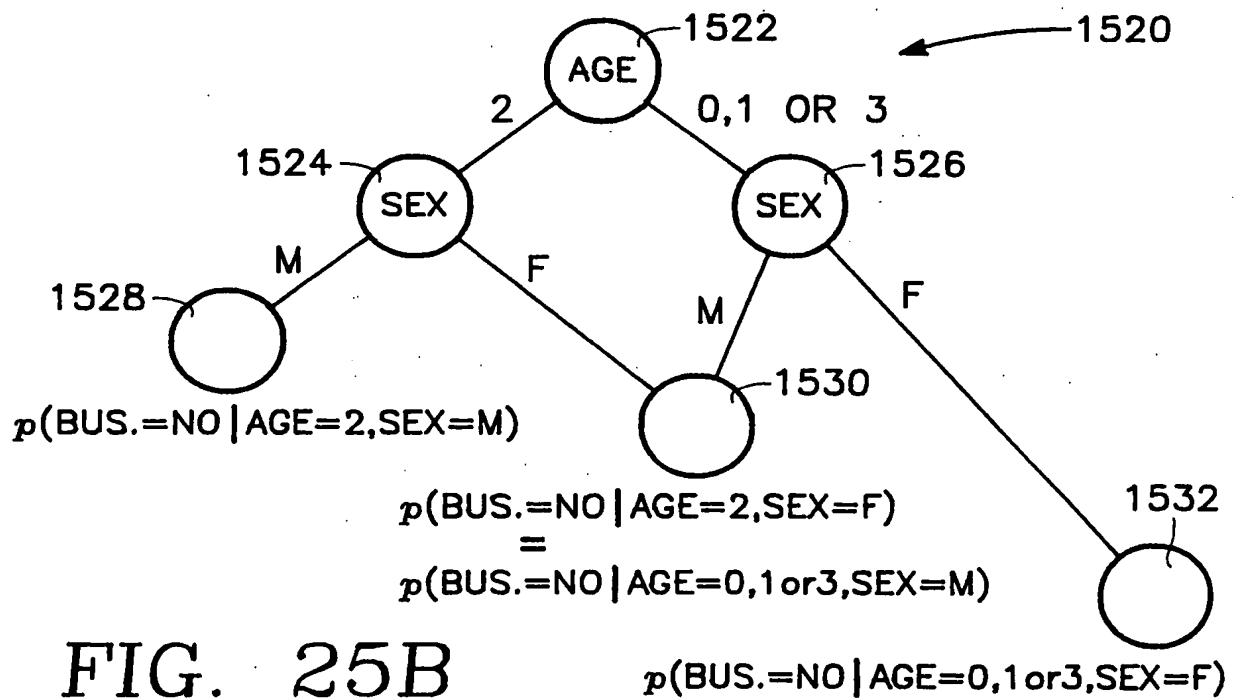



FIG. 22





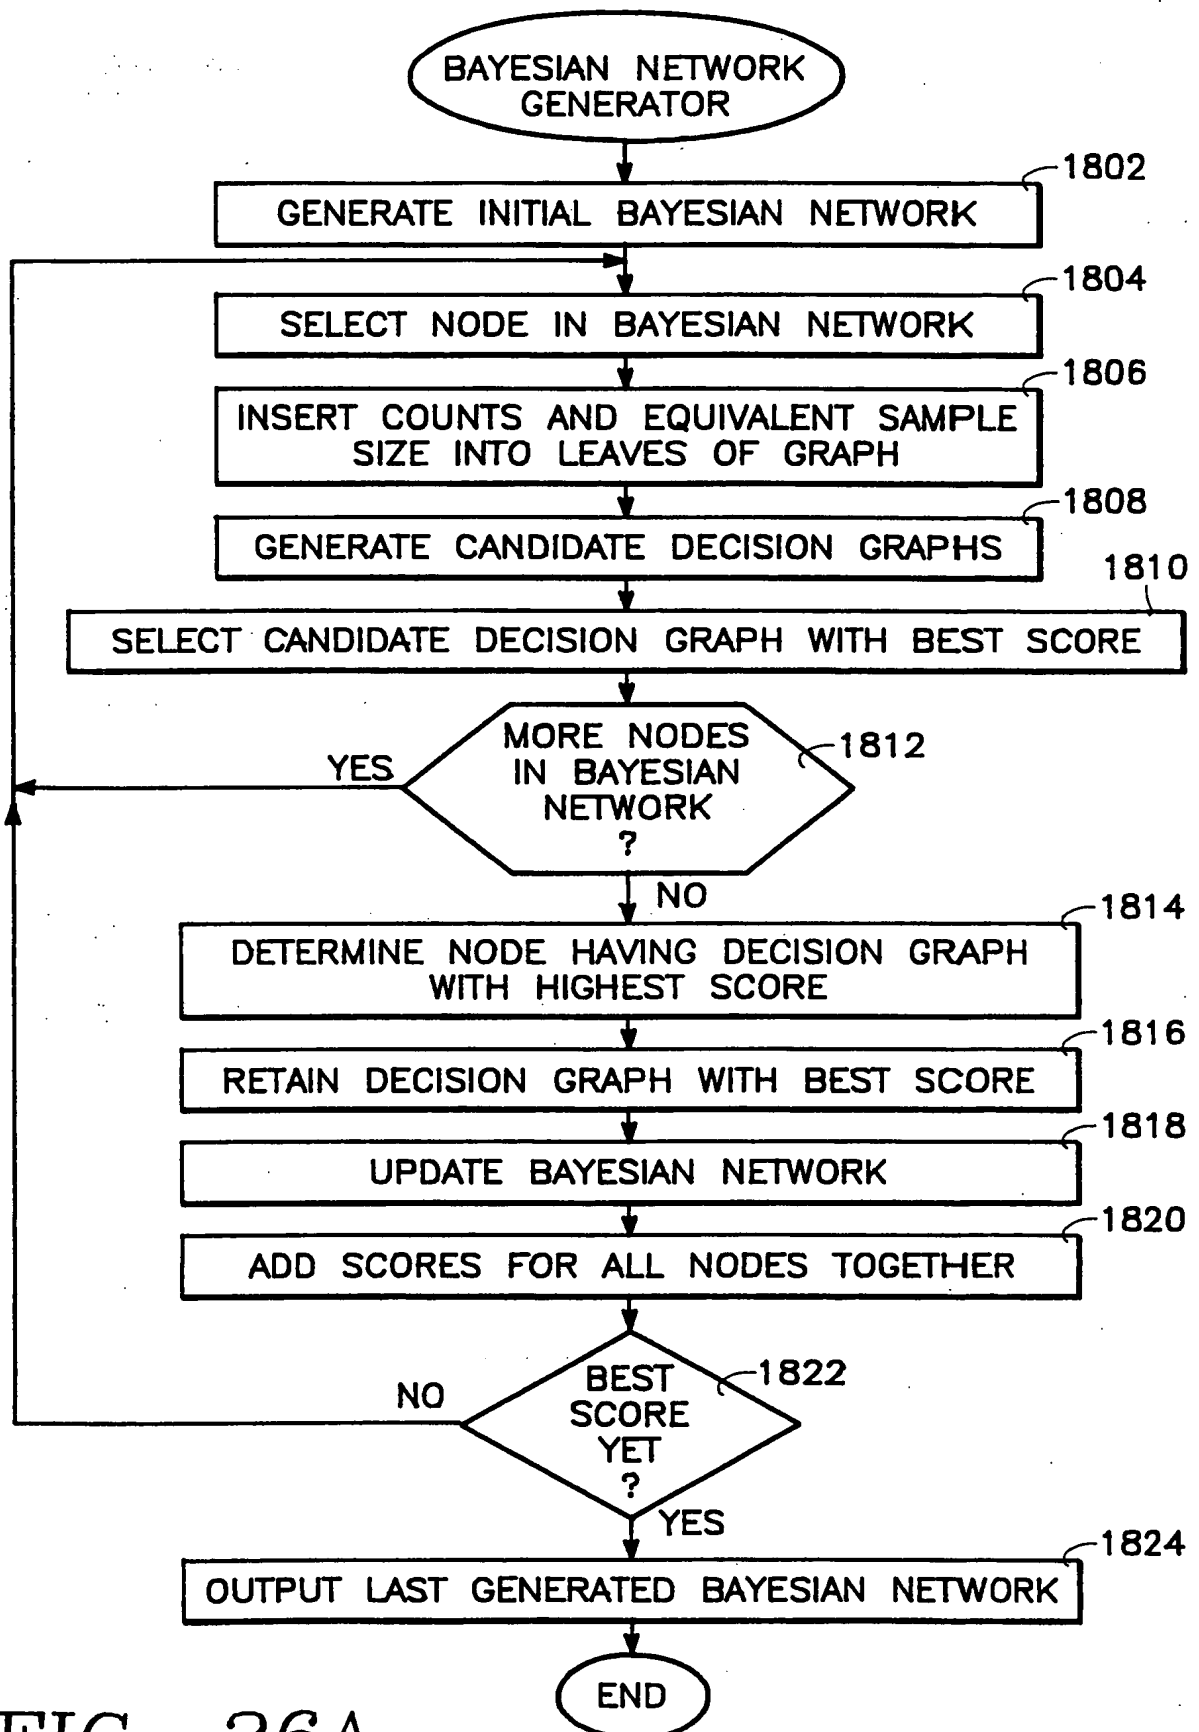


FIG. 26A

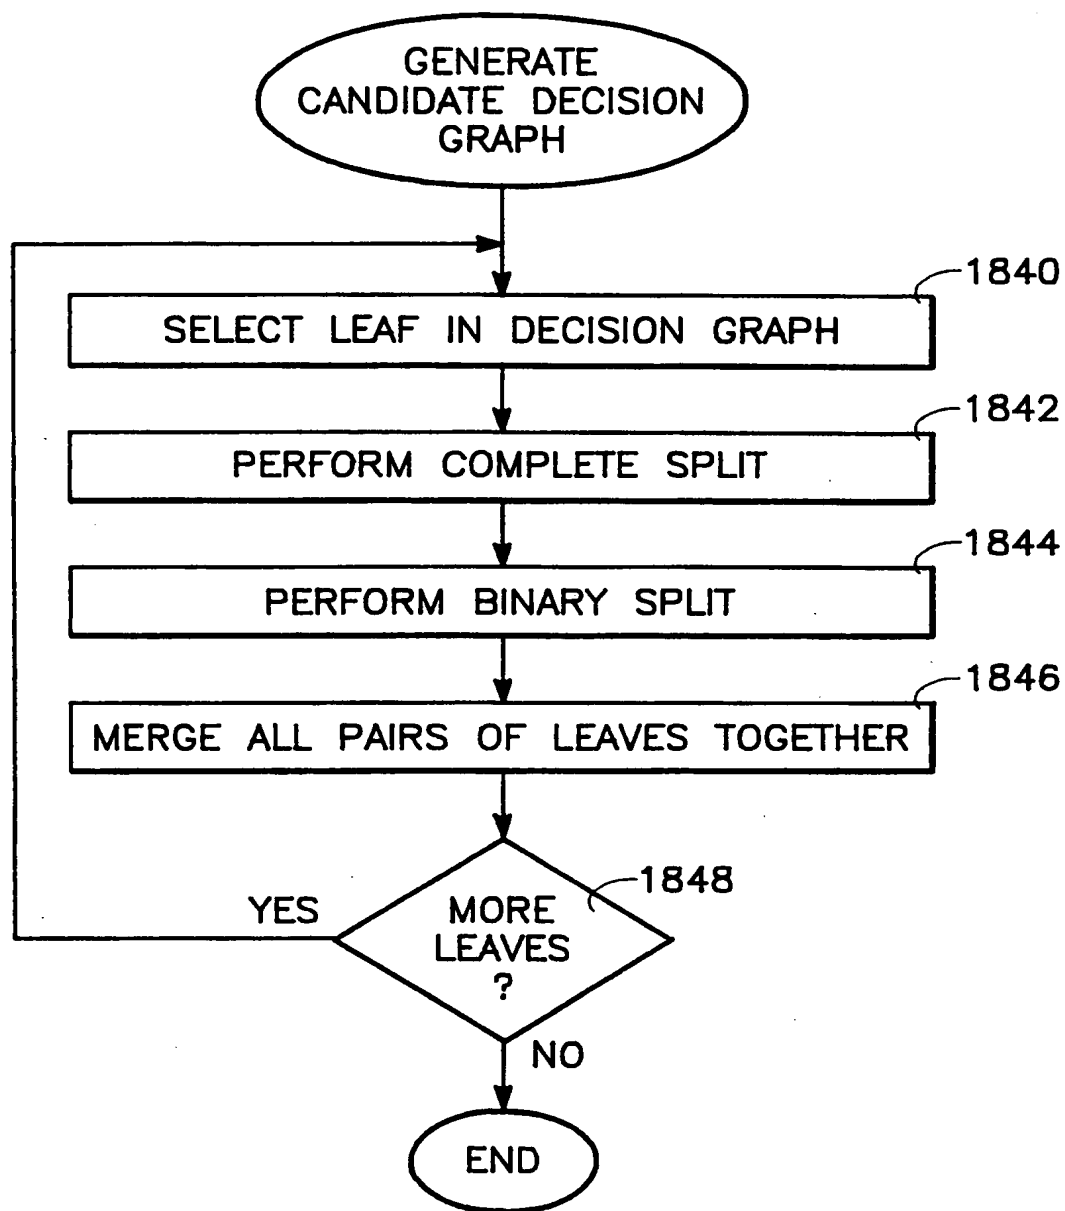


FIG. 26B

FIG. 27A





FIG. 27B




FIG. 27C




FIG. 27D



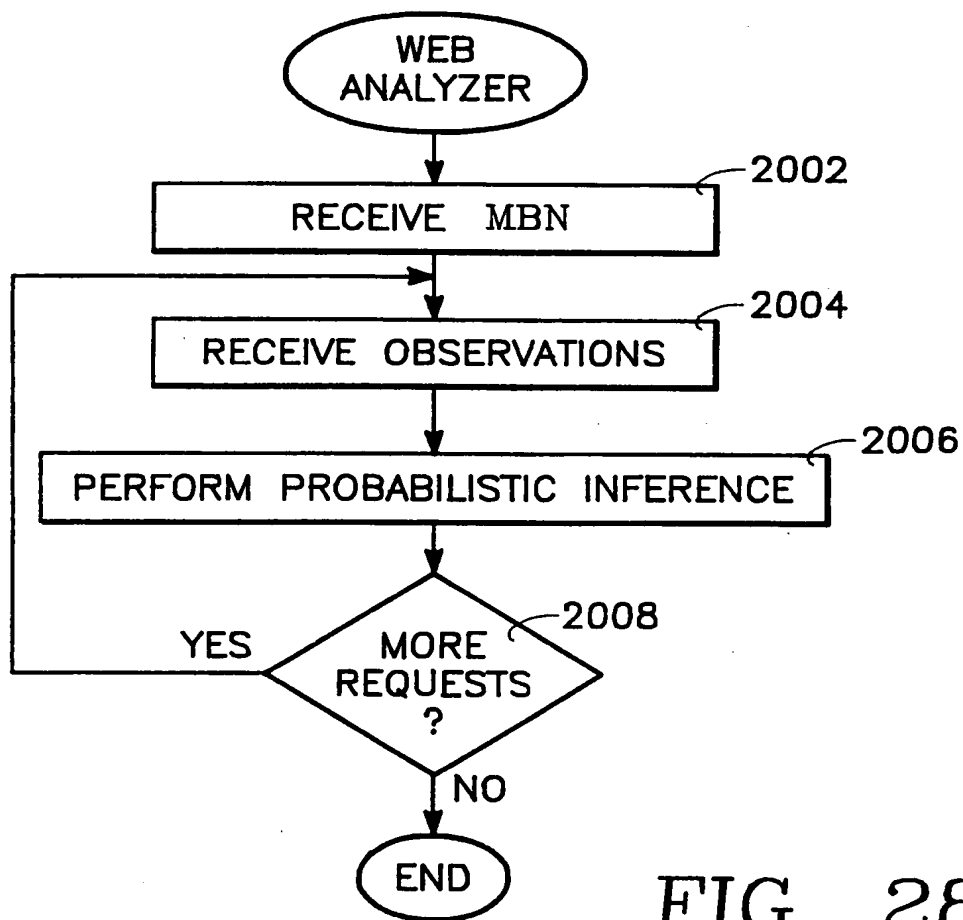


FIG. 28

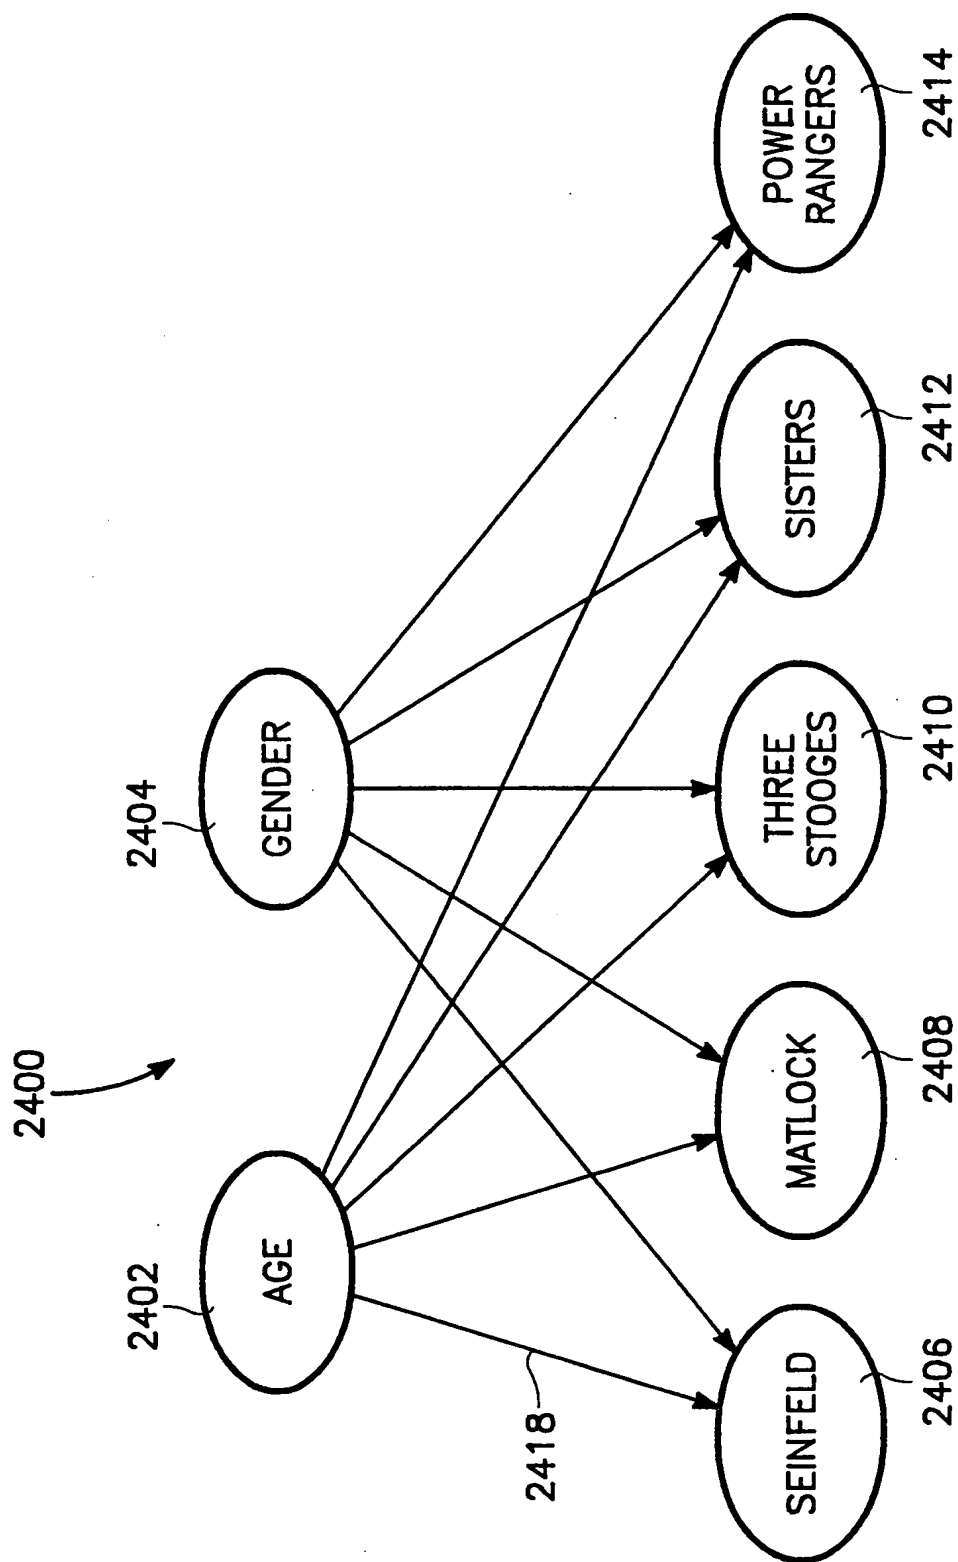


FIG. 29

**PCT**WORLD INTELLECTUAL PROPERTY ORGANIZATION  
International Bureau

## INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<b>(51) International Patent Classification <sup>6</sup> :</b> <b>G06F 9/44</b>	<b>A3</b>	<b>(11) International Publication Number:</b> <b>WO 99/28832</b> <b>(43) International Publication Date:</b> 10 June 1999 (10.06.99)
<b>(21) International Application Number:</b> PCT/US98/25535 <b>(22) International Filing Date:</b> 3 December 1998 (03.12.98)  <b>(30) Priority Data:</b> 09/985,114 4 December 1997 (04.12.97) US  <b>(71) Applicant:</b> MICROSOFT CORPORATION [US/US]; One Microsoft Way, Redmond, WA 98052 (US).  <b>(72) Inventors:</b> THIESSON, Bo; 1502 3rd Place, Kirkland, WA 98033 (US). MEEK, Christopher, A.; 12935 N.E. 71st Street, Kirkland, WA 98033 (US). CHICKERING, David, Maxwell; 7573 Old Redmond Road #12, Redmond, WA 98052 (US). HECKERMAN, David, Earl; 648 West LK Sammamish Lane N.E., Bellevue, WA 98008 (US).  <b>(74) Agent:</b> MICHAELSON, Peter, L.; Michaelson & Wallace, Parkway 109 Office Center, 328 Newman Springs Road, P.O. Box 8489, Red Bank, NJ 07701 (US).		<b>(81) Designated States:</b> CA, CN, JP, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).  <b>Published</b> <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>  <b>(88) Date of publication of the international search report:</b> 22 July 1999 (22.07.99)
<b>(54) Title:</b> MIXTURES OF BAYESIAN NETWORKS  <b>(57) Abstract</b>  One aspect of the invention is the construction of mixtures of Bayesian networks. Another aspect of the invention is the use of such mixtures of Bayesian networks to perform inferencing. A mixture of Bayesian networks (MBN) consists of plural hypothesis-specific Bayesian networks (HSBNs) having possibly hidden and observed variables. A common external hidden variable is associated with the MBN, but is not included in any of the HSBNS. The number of HSBNS in the MBN corresponds to the number of states of the common external hidden variable, and each HSBNS is based upon the hypothesis that the common external hidden variable is in a corresponding one of those states. In one mode of the invention, the MBN having the highest MBN score is selected for use in performing inferencing. In another mode of the invention, some or all of the MBNs are retained as a collection of MBNs which perform inferencing in parallel, their outputs being weighted in accordance with the corresponding MBN scores and the MBN collection output being the weighted sum of all the MBN outputs. In one application of the invention, collaborative filtering may be performed by defining the observed variables to be choices made among a sample of users and the hidden variables to be the preferences of those users.		

**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

# INTERNATIONAL SEARCH REPORT

Inte. .onal Application No

PCT/US 98/25535

**A. CLASSIFICATION OF SUBJECT MATTER**  
IPC 6 G06F9/44

According to International Patent Classification (IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)  
IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>DAVID HECKERMAN: "A Tutorial on Learning with Bayesian Networks"</p> <p>TECHNICAL REPORT MSR-TR-95-06, November 1996, pages 1-57, XP002103817</p> <p>Microsoft Corporation, Redmond, USA</p> <p>* Retrieved from Internet</p> <p>&lt;URL:ftp://ftp.research.microsoft.com/pub/tr/tr-95-06.ps&gt; 26 May 1999 *</p> <p>see page 1, line 1 - page 20, line 18</p> <p>see page 22, line 16 - page 27, line 6</p> <p>see page 30, line 7 - page 35, line 24</p> <p>see page 37, line 11 - page 39, line 7;</p> <p>figure 7</p> <p style="text-align: center;">-----</p>	1-84

☐ Further documents are listed in the continuation of box C.

☐ Patent family members are listed in annex.

\* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

26 May 1999

Date of mailing of the international search report

08/06/1999

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,  
Fax: (+31-70) 340-3016

Authorized officer

Kingma, Y